arXiv:1303.3233v1 [cs.DB] 13 Mar 2013

Consistency Checking and Querying in Probabilistic Databases under Integrity Constraints

Sergio Flesca[a], Filippo Furfaro[a], Francesco Parisi[a]

[a]*DIMES, University of Calabria, Via Bucci - Rende (CS), Italy*

Abstract

We address the issue of incorporating a particular yet expressive form of integrity constraints (namely, denial constraints) into probabilistic databases. To this aim, we move away from the common way of giving semantics to probabilistic databases, which relies on considering a unique interpretation of the data, and address two fundamental problems: *consistency checking* and *query evaluation*. The former consists in verifying whether there is an interpretation which conforms to both the marginal probabilities of the tuples and the integrity constraints. The latter is the problem of answering queries under a "cautious" paradigm, taking into account all interpretations of the data in accordance with the constraints. In this setting, we investigate the complexity of the above-mentioned problems, and identify several tractable cases of practical relevance.

Keywords: Probabilistic databases, Integrity constraints, Consistency checking

1. Introduction

Probabilistic databases (PDBs) are widely used to represent uncertain information in several contexts, ranging from data collected from sensor networks, data integration from heterogeneous sources, bio-medical data, and, more in general, data resulting from statistical analyses. In this setting, several relevant results have been obtained regarding the evaluation of conjunctive queries, thanks to the definition of probabilistic frameworks dealing with two substantially different scenarios: the case of *tuple-independent* PDBs [11, 24], where all the tuples of the database are considered independent one from another, and the case of PDBs representing probabilistic networks encoding even complex forms of correlations among the data [44]. However, none of these frameworks takes into account integrity constraints in the same way as it happens in the deterministic setting, where constraints are used to enforce the consistency of the data. In fact, the former framework strongly relies on the independence assumption (which clearly is in contrast with the presence of the correlations entailed by integrity constraints). The latter framework is closer to an AI perspective of representing the information, as it requires the correlations among the data to be represented as data themselves. This is different from the DB perspective, where constraints are part of the schema, and not of the data.

In this paper, we address the issue of incorporating integrity constraints into probabilistic databases, with the aim of extending the classical semantics and usage of integrity constraints of the deterministic setting to the probabilistic one. Specifically, we consider one of the most popular logical models for the probabilistic data, where information is represented into tuples associated with probabilities, and give the possibility of imposing *denial constraints* on the data, i.e., constraints forbidding the co-existence of certain tuples. In our framework, the role of integrity constraints

Email addresses: flesca@dimes.unical.it (Sergio Flesca), furfaro@dimes.unical.it (Filippo Furfaro), fparisi@dimes.unical.it (Francesco Parisi)

is the same as in the deterministic setting: they can be used to decide whether a new tuple can be inserted in the database, or to decide (a posteriori w.r.t. the generation of the data) if the data are consistent.

Before explaining in detail the main contribution of our work, we provide a motivating example, which clarifies the impact of augmenting a PDB with (denial) constraints. In particular, we focus on the implications on the consistency of the probabilistic data, and on the evaluation of queries. We assume that the reader is acquainted with the data representation model where uncertainty is represented by associating tuples with a probability, and with the notion of possible world. (however, these concepts will be formally recalled in the first sections of the paper).

Motivating Example

Consider the PDB schema $\mathcal{D}^p$ consisting of the relation schema *Room^p(Id, Hid, Price, Type, View, P)*, and its instance *room^p* in Figure 1.

	Id	***Hid***	***Price***	***Type***	***View***	***P***
t_1	1	1	120	Std	Sea	p_1
t_2	2	1	70	Suite	Courtyard	p_2
t_3	3	1	120	Std	Sea	p_3

Figure 1. Relation instance *room^p*

Every tuple in *room^p* is characterized by the room identifier *Id*, the identifier *Hid* of the hotel owning the room, its price per night, its type (e.g., "Standard", "Suite"), and the attribute *View* describing the room view. The attribute *P* specifies the probability that the tuple is true. For now, we leave the probabilities of the three tuples as parameters (p_1, p_2, p_3), as we will consider different values to better explain the main issues related to the consistency and the query evaluation.

Assume that the following constraint *ic* is defined over $\mathcal{D}^p$: "*in the same hotel, standard rooms cannot be more expensive than suites*". This is a denial constraint, as it forbids the coexistence of tuples not satisfying the specified property. In particular, *ic* entails that t_1 and t_2 are mutually exclusive, as, according to t_1, the standard room 1 would be more expensive than the suite room 2 belonging to the same hotel as room 1. For the same reason, *ic* forbids the coexistence of t_2 and t_3.

Finally, consider the following query q on $\mathcal{D}^p$: "*Are there two standard rooms with sea view in hotel* 1*?*". We now show how the consistency of the database and the answer of q vary when changing the probabilities of *room^p*'s tuples.

Case 1 (No admissible interpretation): $p_1 = \frac{3}{4}$; $p_2 = \frac{1}{2}$; $p_3 = \frac{1}{2}$.
In this case, we can conclude that the database is inconsistent. In fact, *ic* forbids the coexistence of t_1 and t_2, which means that the possible worlds containing t_1 must be distinct from those containing t_2. But the marginal probabilities of t_1 and t_2 do not allow this: the fact that $p_1 = \frac{3}{4}$ and $p_2 = \frac{1}{2}$ implies that the sum of the probabilities of the worlds containing either t_1 or t_2 would be $\frac{3}{4}+\frac{1}{2}$, which is greater than 1.

Case 2 (Unique admissible interpretation): $p_1 = \frac{1}{2}$; $p_2 = \frac{1}{2}$; $p_3 = \frac{1}{2}$.
In this case, the database is consistent, as it represents two possible worlds: $w_1 = \{t_1, t_3\}$ and $w_2 = \{t_2\}$, both with probability $\frac{1}{2}$ (correspondingly, the possible worlds representing the other subsets of $\{t_1, t_2, t_3\}$ have probability 0). Observe that there is no other way to interpret the database, while making the constraint satisfied in each possible world, and the probabilities of the possible worlds compatible w.r.t. the marginal probabilities of t_1, t_2, t_3. Thus, the database is consistent and has a unique admissible interpretation.
Now, evaluating the above-defined query q over all the admissible interpretations of the database yields the answer *true* with probability $\frac{1}{2}$ (which is the probability of w_1, the only non-zero-probability world, in the unique admissible interpretation, where q evaluates to *true*). Note that, if *ic* were disregarded and q were evaluated using the independence assumption, the answer of q would be *true* with probability $\frac{1}{4}$.

Case 3 (Multiple admissible interpretations): $p_1 = \frac{1}{2}$; $p_2 = \frac{1}{4}$; $p_3 = \frac{1}{2}$.
In this case, we can conclude that the database is consistent, as it admits at least the interpretations I_1 and I_2 represented in the two rows of the following table (each cell is the probability of the possible world reported in the column header).

	$\emptyset$	$\{t_1\}$	$\{t_2\}$	$\{t_3\}$	$\{t_1, t_2\}$	$\{t_1, t_3\}$	$\{t_2, t_3\}$	$\{t_1, t_2, t_3\}$
I_1	0	1/4	1/4	1/4	0	1/4	0	0
I_2	1/4	0	1/4	0	0	1/2	0	0

With a little effort, the reader can check that there are infinitely many ways of interpreting the database while satisfying the constraints: each interpretation can be obtained by assigning to the possible world $\{t_1, t_3\}$ a different probability in the range $[\frac{1}{4}, \frac{1}{2}]$, and then suitably modifying the probabilities of the other possible worlds where *ic* is satisfied. Basically, the interpretations I_1 and I_2 correspond to the two extreme possible scenarios where, compatibly with the integrity constraint *ic*, a strong negative or positive correlation exists between t_1 and t_3. The other interpretations correspond to scenarios where an "intermediate" correlation exists between t_1 and t_3. Thus, differently from the previous case, there is now more than one admissible interpretation for the database.
Observe that, in the absence of any additional information about the actual correlation among the tuples of $room^p$, all of the above-described admissible interpretations are equally reasonable. Hence, when evaluating queries, we use a "cautious" paradigm, where all the admissible interpretations are taken into account – meaning that no assumption on the actual correlations among tuples is made, besides those which are derivable from the integrity constraints. Thus, according to this paradigm, the answer of query q is *true* with a probability range $[\frac{1}{4}, \frac{1}{2}]$ (where the boundaries of this range are the overall probabilities assigned to the possible worlds containing both t_1 and t_3 by I_1 and I_2). As pointed out in the discussion of Case 2, if the independence assumption were adopted (and *ic* disregarded), the answer of q would be *true* with probability $\frac{1}{4}$, which is the left boundary of the probability range got as cautious answer.

Main contribution

We address the following two fundamental problems:

1) *Consistency checking*: the problem of deciding the consistency of a PDB w.r.t. a given set of denial constraints, that is deciding if there is at least one admissible interpretation of the data. This problem naturally arises when integrity constraints are considered over PDBs: the information encoded in the data (which are typically uncertain) may be in contrast with the information encoded in the constraints (which are typically certain, as they express well-established knowledge about the data domain). Hence, detecting possible inconsistencies arising from the co-existence of certain and uncertain information is relevant in several contexts, such as query evaluation, data cleaning and repairing.
 In this regard, our contribution consists in a thorough characterization of the complexity of this problem. Specifically, after noticing that, in the general case, this problem is *NP*-complete (owing to its interconnection to the probabilistic version of SAT), we identify several islands of tractability, which hold when either:
 i) the conflict hypergraph (i.e., the hypergraph whose edges are the sets of tuples which can not coexist according to the constraints) has some structural property (namely, it is a hypertree or a ring), or
 ii) the constraints have some syntactic properties (independently from the shape of the conflict hypergraph).
2) *Query evaluation*: the problem of evaluating queries over a database which is consistent w.r.t. a given set of denial constraints. Query evaluation relies on the "cautious" paradigm described in Case 3 of the motivating example above, which takes into account all the possible ways of interpreting the data in accordance with the constraints. Specifically, query answers consist of pairs $\langle t, r_p \rangle$, where t is a tuple and r_p a range of probabilities. Therein, r_p is the narrowest interval containing all the probabilities which would be obtained for t as an answer of the query when considering all the admissible interpretations of the data (and, thus, all the correlations among the data compatible with the constraints).
 For this problem, we address both its decisional and search versions, studying the sensitivity of their complexity to the specific constraints imposed on the data and the characteristics of the query. We show that, in the case of general conjunctive queries, the query evaluation problem is $FP^{NP[\log n]}$-hard and in FP^{NP} (note that FP^{NP} is contained in #*P*, the class for which the query evaluation problem under the independence assumption is complete). Moreover, we identify tractable cases where the query evaluation problem is in *PTIME*, which depend on the characteristics of the query and, analogously to the case of the consistency checking problem, on either the syntactic form of the constraints or on some structural properties of the conflict hypergraph.

Moreover, we consider the following extensions of the framework and discuss their impact on the above-summarized results:

A) *tuples are associated with probability ranges, rather than single probabilities*: this is useful when the data acquisition process is not able to assign a precise probability value to the tuples [31, 35];
B) *also denial constraints are probabilistic*: this allows also the domain knowledge encoded by the constraints to be taken into account as uncertain;
C) *pairs of tuples are considered independent unless this contradicts the constraints*: this is a way of interpreting the data in between adopting tuple-independence and rejecting it, and is well suited for those cases where one finds it reasonable to assume some groups of tuples as independent from one another. For instance, if we consider further tuples pertaining to a different hotel in the introductory example (where constraints involve tuples over the same hotel), it may be reasonable to assume that these tuples encode events independent from those pertaining hotel 1.

2. Fundamental notions

2.1. Deterministic Databases and Constraints

We assume classical notions of database schema, relation schema, and relation instance. Relation schemas will be represented by sorted predicates of the form $R(A_1, \ldots, A_n)$, where R is said to be the name of the relation schema and $A_1, \ldots, A_n$ are attribute names composing the set denoted as $Attr(R)$. A tuple over a relation schema $R(A_1, \ldots, A_n)$ is a member of $\Delta_1 \times \cdots \times \Delta_n$, where each Δ_i is the domain of attribute A_i (with $i \in [1..n]$). A relation instance of R is a set r of tuples over R. A database schema $\mathcal{D}$ is a set of relation schemas, and a database instance D of $\mathcal{D}$ is a set of relation instances of the relation schemas of $\mathcal{D}$. Given a tuple t, the value of attribute A of t will be denoted as $t[A]$.

A *denial constraint* over a database schema $\mathcal{D}$ is of the form $\forall \vec{x}. \neg[R_1(\vec{x}_1) \wedge \cdots \wedge R_m(\vec{x}_m) \wedge \phi(\vec{x})]$, where:

- $R_1, \ldots, R_m$ are name of relations in $\mathcal{D}$;
- $\vec{x}$ is a tuple of variables and $\vec{x}_1, \ldots, \vec{x}_m$ are tuples of variables and constants such that $\vec{x} = \mathrm{Var}(\vec{x}_1) \cup \cdots \cup \mathrm{Var}(\vec{x}_m)$, where $\mathrm{Var}(\vec{x}_i)$ denotes the set of variables in $\vec{x}_i$;
- $\phi(\vec{x})$ is a conjunction of built-in predicates of the form $x \diamond y$, where x and y are either variables in $\vec{x}$ or constants, and $\diamond$ is a comparison operator in $\{=, \neq, \leq, \geq, <, >\}$.

m is said to be the *arity* of the constraint. Denial constraints of arity 2 are said to be *binary*. For the sake of brevity, constraints will be written in the form: $\neg[R_1(\vec{x}_1) \wedge \cdots \wedge R_m(\vec{x}_m) \wedge \phi(\vec{x})]$, thus omitting the quantification $\forall \vec{x}$.

We say that a denial constraint ic is *join-free* if no variable occurs in two distinct relation atoms of ic, and, for each built-in predicate occurring in ϕ, at least one term is a constant. Observe that join-free constraints allow multiple occurrences of the same relation name.

It is worth noting that denial constraints enable *equality generating dependencies* (EGDs) to be expressed: an EGD is a denial constraint where all the conjuncts of ϕ are not-equal predicates. Obviously, this means that a denial constraints enables also a functional dependency (FD) to be expressed, as an FD is a binary EGDs over a unique relation (when referring to FDs, we consider also non-canonical ones, i.e., FDs whose RHSs contain multiple attributes).

Given an instance D of the database schema $\mathcal{D}$ and an integrity constraint ic over $\mathcal{D}$, the fact that D satisfies (resp., does not satify) ic is denoted as $D \models ic$ (resp., $D \not\models ic$) and is defined in the standard way. D is said to be consistent w.r.t. a set of integrity constraints $\mathcal{IC}$, denoted with $D \models \mathcal{IC}$, iff $\forall ic \in \mathcal{IC}\ D \models ic$.

Example 1. *Let $\mathcal{D}$ be the (deterministic) database schema consisting of the relation schema* Room(Id, Hid, Price, Type, View), *obtained by removing the probability attribute from the relation schema of our motivating example. Assume the following denial constraints over $\mathcal{D}$:*

ic: $\neg[\mathrm{Room}(x_1, x_2, x_3, \text{'Std'}, x_4) \wedge \mathrm{Room}(x_5, x_2, x_6, \text{'Suite'}, x_7) \wedge\ x_3 > x_6]$, *saying that, in the same hotel, there can not be standard rooms more expensive than suites;*

ic′: $\neg[\mathrm{Room}(x_1, x_2, x_3, x_4, x_5) \wedge \mathrm{Room}(x_6, x_2, x_7, x_4, x_8) \wedge x_3 \neq x_7]$, *imposing that rooms of the same type and hotel have the same price. Thus, ic′ is the FD:* HId, Type→ Price.

where ic is the constraint presented in the introductory example. Consider the relation instance room *of* Room, *obtained from the instance* room^p *of the motivating example by removing column* P. *It is easy to see that* room *satisfies ic′, but does not satisfy ic, since, for the same hotel, the price of standard rooms (rooms* 1 *and* 3*) is greater than that of suite room* 2. □

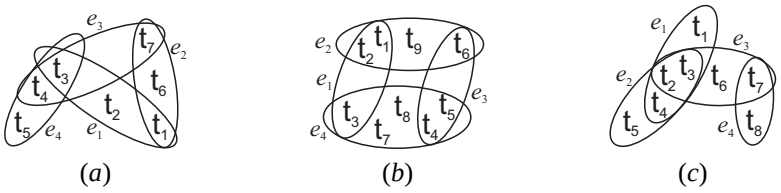


Figure 2. An example of hypergraph (a), ring (b), hypertree (c)

2.2. Hypergraphs and hypertrees

A hypergraph is a pair $H = \langle N, E \rangle$, where N is a set of *nodes*, and E a set of subsets of N, called the *hyperedges* of H. The sets N and E will be also denoted as $N(H)$ and $E(H)$, respectively. Hypergraphs generalize graphs, as graphs are hypergraphs whose hyperedges have exactly two elements (and are called *edges*). Examples of hypergraphs are depicted in Figure 2.

Given a hypergraph $H = \langle N, E \rangle$ and a pair of its nodes n_1, n_2, a *path* connecting n_1 and n_2 is a sequence $e_1, \ldots, e_m$ of distinct hyperedges of H (with $m \geq 1$) such that $n_1 \in e_1$, $n_2 \in e_m$ and, for each $i \in [1..m-1]$, $e_i \cap e_{i+1} \neq \emptyset$. A path connecting n_1 and n_2 is said to be *trivial* if $m = 1$, that is, if it consists of a single edge containing both nodes.

Let $\mathcal{R} = e_1, \ldots, e_m$ be a sequence of hyperedges. We say that e_i and e_j are *neighbors* if $j = i + 1$, or $i = m$ and $j = 1$ (or: if $i = j + 1$, or $i = 1$ and $j = m$). The sequence $\mathcal{R}$ is said to be a *ring* if: *i*) $m \geq 3$; *ii*) for each pair e_i, e_j ($i \neq j$), it holds that $e_i \cap e_j \neq \emptyset$ if and only if e_i and e_j are neighbors. An example of ring is depicted in Figure 2(b). It is easy to see that the definition of ring collapses to the definition of cycle in the case that the hypergraph is a graph.

The nodes appearing in a unique edge will be said to be *ears* of that edge. The set of ears of an edge e will be denoted as $ears(e)$. For instance, in Figure 2(a), $ears(e_1) = \{t_2\}$ and $ears(e_3) = \emptyset$.

A set of nodes N' of H is said to be an *edge-equivalent set* if all the nodes in N' appear altogether in the edges of H. That is, for each $e \in E$ such that $e \cap N' \neq \emptyset$, it holds that $e \cap N' = N'$. Equivalently, the nodes in N' are said to be *edge-equivalent*. For instance, in the hypergraph of Figure 2(b), $\{t_1, t_2\}$ is an edge-equivalent set, as both t_1 and t_2 belong to the edges e_1, e_2 only. Analogously, in the hypergraph of Figure 2(c), nodes t_2 and t_3 are edge equivalent, while $\{t_2, t_3, t_4\}$ is not an edge-equivalent set. Observe that sets of nodes which do not belong to any edge, as well as the ears of an edge (which belong to one edge only), are particular cases of edge-equivalent sets.

A hypergraph is said to be *connected* iff, for each pair of its nodes, a path connects them. A hypergraph H is a *hypertree* iff it is connected and it satisfies the following *acyclicity* property: there is no pair of edges e_1, e_2 such that removing the nodes composing their intersection from every edge of H results in a new hypergraph where the remaining nodes of e_1 are still connected to the remaining nodes of e_2. An instance of hypertree is depicted in Figure 2(c). Observe that the hypergraph in Figure 2(a) is not a hypertree, as the nodes t_2 and t_6 of e_1 and e_2, respectively, are still connected (through the path e_1, e_3, e_2) even if we remove node t_1, which is shared by e_1 and e_2. It is easy to see that hypertrees generalize trees. Basically, the acyclicity property of hypertrees used in this paper is the well-known γ-acyclicity property introduced in [16]. In [15, 16], polynomial time algorithms for checking that a hypergraph is γ-acyclic (and thus a hypertree) are provided.

3. PDBs under integrity constraints

3.1. Probabilistic Databases (PDBs)

A probabilistic relation schema is a classical relation schema with a distinguished attribute P, called *probability*, whose domain is the real interval $[0, 1]$ and which functionally depends on the set of the other attributes. Hence, a probabilistic relation schema has the form $R^p(A_1, \ldots, A_n, P)$. A PDB schema $\mathcal{D}^p$ is a set of probabilistic relation schemas. A probabilistic relation instance r^p is an instance of R^p and a PDB instance D^p is an instance of $\mathcal{D}^p$. We use the superscript p to denote probabilistic relation and database schemas, and their instances. For a tuple $t \in D^p$, the value $t[P]$ is the probability that t belongs to the real world. We also denote $t[P]$ as $p(t)$.

Given a probabilistic relation schema R^p (resp., relation instance r^p, probabilistic tuple t), we write $det(R^p)$ (resp., $det(r^p)$, $det(t)$) to denote its "deterministic" part. Hence, given $R^p(A_1, \ldots, A_n, P)$, $det(R^p) = R(A_1, \ldots, A_n)$, and $det(r^p) = \pi_{Attr(det(R^p))}(r^p)$, and $det(t) = \pi_{Attr(det(R^p))}(t)$. This definition is extended to deal with the deterministic part of PDB schemas and instances in the obvious way.

3.1.1. Possible world semantics

The semantics of a PDB is based on *possible worlds*. Given a PDB D^p, a possible world is any subset of its deterministic part $det(D^p)$. The set of possible worlds of D^p is as follows: $pwd(D^p) = \{w \mid w \subseteq det(D^p)\}$. An Pr interpretation of D^p is a *probability distribution function* (PDF) over the set of possible worlds $pwd(D^p)$ which satisfies the following property:

$$(i)\ \ \forall t \in D^p,\ \ p(t) = \sum_{\substack{w \in pwd(D^p) \\ \wedge\ det(t) \in w}} Pr(w).$$

Condition (i) imposes that the probability of each tuple t of D^p coincides with that specified in t itself. Observe that, from definition of PDF, Pr must also satisfy the following conditions:

$$(ii) \sum_{w \in pwd(D^p)} Pr(w) = 1; \qquad (iii)\ \ \forall w \in pwd(D^p),\ Pr(w) \geq 0;$$

meaning that Pr assigns a non-negative probability to each possible world, and that the probabilities assigned by Pr to the possible worlds sum up to 1.

The set of interpretations of a PDB D^p will be denoted as $\mathcal{I}(D^p)$.

Observe that, strictly speaking, possible worlds are sets of deterministic counterparts of probabilistic tuples. However, for the sake of simplicity, with a little abuse of notation, in the following we will say that a probabilistic tuple t belongs (resp., does not belong) to a possible world w – written $t \in w$ (resp., $t \notin w$) – if w contains (resp., does not contain) the deterministic counterpart of t, i.e., $det(t) \in w$ (resp., $det(t) \notin w$). Moreover, given a deterministic tuple t, we will write $p(t)$ to denote the probability associated with the probabilistic counterpart of t. Thus, $p(t)$ will denote either $t[P]$, in the case that t is a probabilistic tuple, or $t'[P]$, in the case that t is deterministic and t' is its probabilistic counterpart.

If independence among tuples is assumed, only one interpretation of D^p is considered, assigning to each possible world w the probability $Pr(w) = \prod_{t \in w} p(t) \times \prod_{t \notin w}(1-p(t))$. In fact, under the independence assumption, the probability of a conjunct of events is equal to the product of their probabilities. In turn, queries over the PDB are evaluated by considering this unique interpretation. In this paper, we consider a different framework, where independence among tuples is not assumed, and all the possible interpretations are considered.

Example 2. *Consider the PDB schema $\mathcal{D}^p$ and its instance D^p introduced in our motivating example. D^p consists of the relation instance* room^p *reported in Figure 1. Assume that t_1, t_2, t_3 have probabilities $p_1 = p_2 = p_3 = 1/2$, and disregard the integrity constraint defined in the motivating example.*

Table 1 shows some interpretations of D^p. Pr_1 corresponds to the interpretation obtained by assuming tuple independence. Interpretation Pr_5, where ϵ is any real number in $[0, 1/4]$, suffices to show that there are infinitely many interpretations of D^p. □

Possible worlds (w)

Interpretations	$\emptyset$	$\{t_1\}$	$\{t_2\}$	$\{t_3\}$	$\{t_1,t_2\}$	$\{t_1,t_3\}$	$\{t_2,t_3\}$	$\{t_1,t_2,t_3\}$	
$Pr_1(w)$	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8	} Assuming tuple independence
$Pr_2(w)$	0	1/2	0	0	0	0	1/2	0	} Further interpretations corresponding to different correlations among tuples
$Pr_3(w)$	0	0	1/2	0	0	1/2	0	0	
$Pr_4(w)$	0	0	0	1/2	1/2	0	0	0	
$Pr_5(w)$	$1/2-2\epsilon$	ϵ	ϵ	ϵ	0	0	0	$1/2-\epsilon$	

Table 1. Some interpretations of D^p

3.2. Imposing denial constraints over PDBs

An integrity constraint over a PDB schema $\mathcal{D}^p$ is written as an integrity constraint over its deterministic part $det(\mathcal{D}^p)$. Its impact on the semantics of the instances of $\mathcal{D}^p$ is as follows. As explained in the previous section, a PDB D^p, instance of $\mathcal{D}^p$, may have several interpretations, all equally sound. However, if some constraints are known

on its schema $\mathcal{D}^p$, some interpretations may have to be rejected. The interpretations to be discarded are those "in contrast" with the domain knowledge expressed by the constraints, that is, those assigning a non-zero probability to worlds violating some constraint.

Formally, given a set of constraints $\mathcal{IC}$ on $\mathcal{D}^p$, an interpretation $Pr \in \mathcal{I}(D^p)$ is *admissible* (and said to be a *model for* D^p *w.r.t.* $\mathcal{IC}$) if $\sum_{w \in pwd(D^p) \wedge w \models \mathcal{IC}} Pr(w) = 1$ (or, equivalently, if $\sum_{w \in pwd(D^p) \wedge w \not\models \mathcal{IC}} Pr(w) = 0$). The set of models of D^p w.r.t. $\mathcal{IC}$ will be denoted as $\mathcal{M}(D^p, \mathcal{IC})$. Obviously, $\mathcal{M}(D^p, \mathcal{IC})$ coincides with the set of interpretations $\mathcal{I}(D^p)$ if no integrity constraint is imposed ($\mathcal{IC} = \emptyset$), while, in general, $\mathcal{M}(D^p, \mathcal{IC}) \subseteq \mathcal{I}(D^p)$.

Example 3. *Consider the PDB D^p and the integrity constraint ic introduced in our motivating example. Assume that all the tuples of* room^p *have probability* $1/2$*. Thus, the interpretations for D^p are those discussed in Example 2 (see also Table 1). It is easy to see that* room^p *admits at least one model, namely Pr_3 (shown in Table 1), which assigns non-zero probability only to $w_1 = \{t_2\}$ and $w_2 = \{t_1, t_3\}$. In fact, it can be proved that Pr_3 is the unique model of* room^p *w.r.t. ic, since every other interpretation of* room^p*, including Pr_1 where tuple independence is assumed, makes the constraint ic violated in some non-zero probability world. This example shows an interesting aspect of denial constraints. Although denial constraints only explicitly forbid the co-existence of tuples, they may implicitly entail the co-existence of tuples: for instance, for the given probabilities of t_1, t_2, t_3, constraint ic implies the coexistence of t_1 and t_3.* □

Example 3 re-examines Case 2 of our motivating example, and shows a case where the PDB is consistent and admits a unique model. The reader is referred to the discussions of Case 1 and Case 3 of the motivating example to consider different scenarios, where the PDB is not consistent (Case 1), or is consistent and admits several models (Case 3).

3.2.1. Modeling denial constraints as hypergraphs

Basically, a denial constraint over a PDB restricts its models w.r.t. the set of interpretations, as it expresses the fact that some sets of tuples of D^p are *conflicting*, that is, they cannot co-exist: an interpretation is not a model if it assigns a non-zero probability to a possible world containing these tuples altogether. Hence, a set of denial constraints $\mathcal{IC}$ can be naturally represented as a *conflict hypergraph*, whose nodes are the tuples of D^p and where each hyperedge consists of a set of tuples whose co-existence is forbidden by a denial constraint in $\mathcal{IC}$ (in fact, hypergraphs were used to model denial constraints also in several works dealing with consistent query answers in the deterministic setting [8]). The definitions of *conflicting tuples* and *conflict hypergraph* are as follows.

Definition 1 (Conflicting set of tuples)**.** *Let $\mathcal{D}^p$ be a PDB schema, $\mathcal{IC}$ a set of denial constraints on $\mathcal{D}^p$, and D^p an instance of $\mathcal{D}^p$. A set T of tuples of D^p is said to be a* conflicting set *w.r.t. $\mathcal{IC}$ if it is a minimal set such that any possible world containing all the tuples in T violates $\mathcal{IC}$.*

Example 4. *In Example 3, both $\{t_1, t_2\}$ and $\{t_2, t_3\}$ are conflicting sets of tuples w.r.t. $\mathcal{IC} = \{ic\}$, while $\{t_1, t_2, t_3\}$ is not, as it is not minimal.* □

Definition 2 (Conflict hypergraph)**.** *Let $\mathcal{D}^p$ be a PDB schema, $\mathcal{IC}$ a set of denial constraints on $\mathcal{D}^p$, and D^p an instance of $\mathcal{D}^p$. The conflict hypergraph of D^p w.r.t. $\mathcal{IC}$ is the hypergraph $HG(D^p, \mathcal{IC})$ whose nodes are the tuples of D^p and whose hyperedges are the conflicting sets of D^p w.r.t. $\mathcal{IC}$.*

Example 5. *Consider a database instance D^p having tuples $t_1, \ldots, t_9$, and a set of denial constraints $\mathcal{IC}$ stating that $e_1 = \{t_1, t_4, t_7\}$, $e_2 = \{t_1, t_2, t_3, t_4, t_5, t_6\}$, $e_3 = \{t_3, t_6, t_9\}$, and $e_4 = \{t_6, t_8\}$ are conflicting sets of tuples. The conflict hypergraph $HG(D^p, \mathcal{IC})$ in Figure 3 concisely represents this fact.* □

It is easy to see that, if $\mathcal{IC}$ contains binary denial constraints only, then the conflict hypergraph collapses to a graph.

Example 6. *Consider D^p and $\mathcal{IC} = \{ic\}$ of our motivating example – observe that ic is a binary denial constraint. The graph representing $HG(D^p, \mathcal{IC})$ is shown in Figure 4.* □

It is easy to see that the size of the conflict hypergraph is polynomial w.r.t. the size of D^p (in particular, its number of nodes is bounded by the number of tuples of D^p) and can be constructed in polynomial time w.r.t. the size of D^p.

Remark 1. Observe that the conflict hypergraph $H(D^p, \mathcal{IC})$ corresponds to a representation of the dual lineage of the *constraint query* $q_{\mathcal{IC}}$, i.e., the boolean query $q_{\mathcal{IC}} = \bigvee_{ic \in \mathcal{IC}}(\neg ic)$ which basically asks whether there is *no* model for D^p

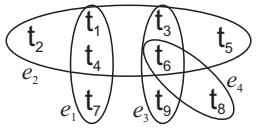


Figure 3. A conflict hypergraph.

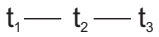


Figure 4. Conflict graph of the motivating example

w.r.t. $\mathcal{IC}$. For instance, consider the case of Example 3. A lineage of $q_{\mathcal{IC}}$ is the DNF expression: $(X_1 \wedge X_2) \vee (X_2 \wedge X_3)$, where each X_i corresponds to tuple t_i. Thus, the semantics of the considered constraints is captured by the dual lineage, that is the CNF expression $(Y_1 \vee Y_2) \wedge (Y_2 \vee Y_3)$, where each $Y_i = \text{not}(X_i)$. It is easy to see that the conflict hypergraph (as described in Example 6) is the hypergraph of this CNF expression. In the conclusions (Section 8), we will elaborate more on this relationship between conflict hypergraphs and (dual) lineages of constraint queries: exploiting this relationship may help to tackle the problems addressed in this paper from a different perspective.

4. Consistency checking

Detecting inconsistencies is fundamental for certifying the quality of the data and extracting reliable information from them. In the deterministic setting, inconsistency typically arises from errors that occurred during the generation of the data, as well as during their acquisition. In the probabilistic setting, there is one more possible source of inconsistency, coming from the technique adopted for estimating the "degree of uncertainty" of the acquired information, which determines the probability values assigned to the probabilistic tuples. Possible bad assignments of probability values can turn out when integrity constraints on the data domain (which typically encode certain information coming from well-established knowledge of the domain) are considered.

In this section, we address the problem of checking this form of consistency, that is, the problem of checking whether the probabilities associated with the tuples are "compatible" with the integrity constraints defined over the data. It is worth noting that the study of this problem has a strong impact in several aspects of the management of probabilistic data: checking the consistency can be used during the data acquisition phase (in order to "certify" the validity of the model applied for determining the probabilities of the tuples), as well as a preliminary step of the computation of the query answers. Moreover, it is strongly interleaved with the problem of repairing the data, whose study is deferred to future work.

Before providing the formal definition of the consistency checking problem, we introduce some basic notions and notations. Given a PDB schema $\mathcal{D}^p$, a set of integrity constraint $\mathcal{IC}$, and an instance D^p of $\mathcal{D}^p$, we say that D^p satisfies (resp., does not satisfy) $\mathcal{IC}$, denoted as $D^p \models \mathcal{IC}$ (resp., $D^p \not\models \mathcal{IC}$) iff the set of models $\mathcal{M}(D^p, \mathcal{IC})$ is not empty. In the following, we will say "*consistent w.r.t.*" (resp., "*inconsistent w.r.t.*") meaning the same as "*satisfies*" (resp., "*does not satisfy*").

We are now ready to provide the formal definition of the consistency checking problem. In this definition, as well as in the rest of the paper, we assume that a PDB schema $\mathcal{D}^p$ and a set of denial constraints $\mathcal{IC}$ over $\mathcal{D}^p$ are given.

Definition 3 (Consistency Checking Problem (cc))**.** *Given a PDB instance D^p of $\mathcal{D}^p$, the consistency checking problem (*cc*) is deciding whether $D^p \models \mathcal{IC}$.*

We point out that, in our complexity analysis, $\mathcal{D}^p$ and $\mathcal{IC}$ will be assumed of fixed size, thus we refer to data complexity.

The following theorem states that cc is NP-complete, and it easily derives from the interconnection of cc with the NP-complete problem PSAT [22] (*Probabilistic satisfiability*), which is the generalization of SAT defined as follows: "*Let $S = \{C_1, \ldots, C_m\}$ be a set of m clauses, where each C_i is a disjunction of literals (i.e, possibly negated propositional variables $x_1, \ldots, x_n$) and each C_i is associated with a probability p_i. Decide whether S is* satisfiable*, that is, whether there is a probability distribution π over all the 2^n possible truth assignments over $x_1, \ldots, x_n$ such that, for each C_i, the sum of the probabilities assigned by π to the truth assignments satisfying C_i is equal to p_i.*" Basically, the membership in NP of cc derives from the fact that any instance of cc over a PDB D^p can be reduced to an equivalent PSAT instance where: *a*) the propositional variables correspond to the tuples of D^p, *b*) the constraints of cc are encoded into clauses with probability 1, *c*) the fact that the tuples are assigned a probability is encoded

into a clause for each tuple, with probability equal to the tuple probability. As regards the hardness of cc for *NP*, it intuitively derives from the fact that the hardness of PSAT was shown in [22] for the case that only unary clauses have probabilities different from 1: thus, this proof can be applied on cc, by mapping unary clauses to tuples and the other clauses (which are deterministic) to constraints [1].

Theorem 1 (Complexity of cc)**.** cc *is NP-complete.*

In the following, we devote our attention to determining tractable cases of cc, from two different perspectives. First, in Section 4.1, we will show tractable cases which depend from the structural properties of the conflict hypergraph, and, thus, from how the data combine with the constraints. The major results of this section are that cc is tractable if the conflict hypergraph is either a hypertree or ring. Then, in Section 4.2, we will show syntactic conditions on the constraints which make cc tractable, independently from the shape of the conflict hypergraph. At the end of the latter section, we also discuss the relationship between these two kinds of tractable cases.

4.1. Tractability arising from the structure of the conflict hypergraph

It is worth noting that, since there is a polynomial-time reduction from cc to PSAT, the tractability results for PSAT may be exploited for devising efficient strategy for solving cc. In fact, in [22], it was shown that 2PSAT (where clauses are binary) can be solved in polynomial time if the graph of clauses (which contains a node for each literal and an edge for each pair of literals occurring in the same clause) is outerplanar. This result relies on a suitable reduction of 2PSAT to a tractable instance of 2MAXSAT (maximum weight satisfiability with at most two literals per clause). Since, in the case of binary denial constraints, the conflict hypergraph is a graph and the above-discussed reduction of cc to PSAT results in an instance of 2PSAT where the graph of clauses has the same "shape" of our conflict graph, we have that cc is polynomial-time solvable if denial constraints are binary and the conflict graph is outerplanar. However, on the whole, reducing 2PSAT to 2MAXSAT and then solving the obtained 2MAXSAT instance require a high polynomial-degree computation (specifically, the complexity is $O(n^6 \log n)$, where n is the number of literals in the PSAT formula, corresponding to the number of tuples in our case).

Here, we detect tractable cases of cc, which, up to our knowledge, are not subsumed by any known tractability result for PSAT. Our tractable cases have the following amenities:

- no limitation is put on the arity of the constraints;
- instead of exploiting reductions of cc to other problems, we determine necessary and sufficient conditions which can be efficiently checked (in linear time) by only examining the conflict hypergraph and the probabilities of the tuples.

Our main results regarding the tractability arising from the structure of the conflict hypergraph (which will be given in sections 4.1.2 and 4.1.3) are that consistency can be checked in linear time over the conflict hypergraph if it is either a hypertree or a ring.

4.1.1. New notations and preliminary results

Before providing our characterization of tractable cases arising from the structure of the conflict hypergraph, we introduce some preliminary results and new notations. Given a hypergraph $H = \langle N, E\rangle$ and a hyperedge $e \in E$, the set of intersections of e with the other hyperedges of H is denoted as $Int(e, H) = \{s \mid \exists e' \in E \text{ s.t. } e' \neq e \wedge s = e \cap e'\}$. For instance, for the hypertree H in Figure 2(c), $Int(e_1, H) = \{\{t_2, t_3\}, \{t_2, t_3, t_4\}\}$. Moreover, given a set of sets S, we call S a *matryoshka* if there is a total ordering $s_1, \ldots, s_n$ of its elements such that, for each $i, j \in [1..n]$ with $i < j$ it holds that $s_1 \subset s_2 \subset \cdots \subset s_n$. For instance, the above-mentioned set $Int(e_1, H)$ is a matryoshka. Finally, given a set of hyperedges S, we denote as H^{-S} the hypergraph obtained from H by removing the edges of S and the nodes in the edges of S which do not belong to any other edge of the remaining part of H. That is, $H^{-S} = \langle N', E'\rangle$, where $E' = E \setminus S$, $N' = \bigcup_{e \in E'} e$. For instance, for the hypergraph H in Figure 2(a), $H^{-\{e_1\}}$ is obtained by removing e_1 from the set of edges of H, and t_2 from the set of its nodes. Analogously, $H^{-\{e_1, e_2\}}$ will not contain edges e_1 and e_2, as well as nodes t_1, t_2, t_6.

[1] However, we will not provide a formal proof of the *NP*-hardness of cc based on this reasoning, that is, based on reducing hard instances of PSAT to cc instances. Indeed, a formal proof of the hardness will be provided for the theorems 5 and 7 introduced in Section 4.2, which are more specific in stating the hardness of cc in that they say that cc is *NP*-hard in the presence of denial constraints of some syntactic forms.

The first preliminary result (Proposition 1) states a general property of hypertrees: any hypertree H contains at least one edge e which is attached to the rest of H so that the set of intersections of e with the other edges of H is a matryoshka. Moreover, removing this edge from H results in a new hypergraph which is still a hypertree. This result is of independent interest, as it allows for reasoning on hypertrees (conforming to the γ-acyclicity property) by using induction on the number of hyperedges: any hypertree with x edges can be viewed as a hypertree with $x-1$ edges which has been augmented with a new edge, attached to the rest of the hypertree by means of sets of nodes encapsulated one to another.

Proposition 1. *Let $H = \langle N, E\rangle$ be a hypertree. Then, there is at least one hyperedge $e \in E$ such that $Int(e, H)$ is a matryoshka. Moreover, $H^{-\{e\}}$ is still a hypertree.*

As an example, consider the hypertree in Figure 2(c). As ensured by Proposition 1, this hypertree contains the edge e_1 whose set of intersections with the other edges is $\{\{t_2, t_3\}, \{t_2, t_3, t_4\}\}$, which is a matryoska. Moreover, removing e_1 from the set of hyperedges, and the ears of e_1 from the set of nodes, still yields a hypertree. The same holds for e_2 and e_4, but not for e_3.

The second preliminary result (which will be stated as Lemma 1) regards the minimum probability that a set of tuples co-exist according to the models of the given PDB. Specifically, given a set of tuples T of the PDB D^p, we denote this minimum probability as $p^{\min}(T)$, whose formal definition is as follows:

$$p^{\min}(T) = \min_{Pr \in \mathcal{M}(D^p, \mathcal{IC})} \left\{ \sum_{w \in pwd(D^p) \wedge T \subseteq w} Pr(w) \right\}$$

The following example clarifies the semantics of $p^{\min}$.

Example 7. *Consider the case discussed in Example 2 (the same as Case 2 of our motivating example, but with $\mathcal{IC} = \emptyset$). Here, every interpretation is a model. Hence, $p^{\min}(t_1, t_3) = 0$, as there is an interpretation (for instance, Pr_2 or Pr_4 in Table 1) which assigns probability 0 to both the possible worlds $\{t_1, t_3\}$ and $\{t_1, t_2, t_3\}$ – the worlds containing both t_1 and t_3. However, if we impose $\mathcal{IC} = \{ic\}$ of the motivating example, we have that $p^{\min}(t_1, t_3) = 1/2$, as according to Pr_3 (the unique model for the database w.r.t. $\mathcal{IC}$) the probabilities of worlds $\{t_1, t_3\}$ and $\{t_1, t_2, t_3\}$ are, respectively, $1/2$ and 0 (hence, their sum is $1/2$).* □

Lemma 1 states that, for any set of tuples $T = \{t_1, \ldots, t_n\}$, independently from how they are connected in the conflict hypergraph, the probability that they co-exist, for every model, has a lower bound which is implied by their marginal probabilities. This lower bound is $\max\left\{0, \sum_{i=1}^{n} p(t_i) - n + 1\right\}$, which is exactly the minimum probability of the co-existence of $t_1, \ldots, t_n$ in two cases: *i*) the case that $t_1, \ldots, t_n$ are pairwise disconnected in the conflict hypergraph (which happens, for instance, in the very special case that $t_1, \ldots, t_n$ are not involved in any constraint); *ii*) the case that the set of intersections of T with the edges of H is a matryoshka. This is interesting, as it depicts a case of tuples correlated through constraints which behave similarly to tuples among which no correlation is expressed by any constraint.

Lemma 1. *Let D^p be an instance of $\mathcal{D}^p$ consistent w.r.t. $\mathcal{IC}$, T a set of tuples of D^p, and let H denote the conflict hypergraph $HG(D^p, \mathcal{IC})$. If either* i) *the tuples in T are pairwise disconnected in H, or* ii) *$Int(T, H)$ is a matryoshka, then $p^{\min}(T) = \max\{0, \sum_{t \in T} p(t) - |T| + 1\}$. Otherwise, this formula provides a lower bound for $p^{min}(T)$.*

4.1.2. Tractability of hypertrees

We are now ready to state our first result on cc tractability.

Theorem 2. *Given an instance D^p of $\mathcal{D}^p$, if $HG(D^p, \mathcal{IC})$ is a hypertree, then $D^p \models \mathcal{IC}$ iff, for each hyperedge e of $HG(D^p, \mathcal{IC})$, it holds that*

$$\sum_{t \in e} p(t) \leq |e| - 1 \tag{1}$$

Proof. ($\Rightarrow$): We first show that if there is a model for D^p w.r.t. $\mathcal{IC}$, then inequality (1) holds for each hyperedge of $HG(D^p, \mathcal{IC})$. Reasoning by contradiction, assume that $D^p \models \mathcal{IC}$ and there is an hyperedge $e = \{t_1, \ldots, t_n\}$ of $HG(D^p, \mathcal{IC})$ such that $\sum_{i=1}^{n} p(t_i) - n + 1 > 0$. Since this value is a lower bound for $p^{min}(t_1, \ldots, t_n)$ (due to Lemma 1),

it holds that every model M for D^p w.r.t. $\mathcal{IC}$ assigns a non-zero probability to some possible world containing all the tuples $t_1, \ldots, t_n$. This contradicts that M is a model, since any possible world containing $t_1, \ldots, t_n$ does not satisfy $\mathcal{IC}$. ($\Leftarrow$): We now prove that if inequality (1) holds for each hyperedge of $HG(D^p, \mathcal{IC})$, then there is a model for D^p w.r.t. $\mathcal{IC}$. We reason by induction on the number of hyperedges of $HG(D^p, \mathcal{IC})$.

The base case is when $HG(D^p, \mathcal{IC})$ consists of a single hyperedge $e = \{t_1, \ldots, t_k\}$. Consider the same database D^p, but impose over it the empty set of denial constraints, instead of $\mathcal{IC}$. Then, from Lemma 1 (case *i*)), we have that there is at least one model M for D^p (w.r.t. the empty set of constraints) such that $\sum_{w \supseteq \{t_1,\ldots,t_k\}} M(w) = \max\left\{0, \sum_{i=1}^{k} p(t_i) - k + 1\right\}$. The term on the right-hand side evaluates to 0, as, from the hypothesis, we have that $\sum_{i=1}^{k} p(t_i) \leq k - 1$. Hence, M is a model for D^p also w.r.t. $\mathcal{IC}$, since the only constraint entailed by $\mathcal{IC}$ is that the tuples $t_1, \ldots, t_k$ can not be altogether in any possible world with non-zero probability.

We now prove the induction step. Consider the case that $H = HG(D^p, \mathcal{IC})$ is a hypertree with n hyperedges. The induction hypothesis is that the property to be shown holds in the presence of any conflict hypergraph consisting of a hypertree with $n - 1$ hyperedges. Let e be a hyperedge of H such that $Int(e, H)$ is a matryoshka, and $H' = H^{-\{e\}}$ is a hypertree. The existence of e and the fact that H' is a hypertree are guaranteed by Proposition 1. We denote the nodes in e as $t'_1, \ldots, t'_m, t''_1, \ldots, t''_n$, where $T' = \{t'_1, \ldots, t'_m\}$ is the set of nodes of e in H', and $T'' = \{t''_1, \ldots t''_n\}$ are the ears of e. Correspondingly, D'' is the portion of D^p containing only the tuples $t''_1, \ldots t''_n$, and D' is the portion of D^p containing all the other tuples (that is, the tuples corresponding to the nodes of H'). We consider D' associated with the set of constraints imposed by H', and D'' associated with an empty set of constraints.

Thanks to the induction hypothesis, and to the fact that inequality (1) holds, we have that D' is consistent w.r.t. the set of constraints encoded by H'. Moreover, since $Int(e, H)$ is a matryoshka, we have that the set T' is such that $Int(T', H')$ is a matrioshka too. Hence, from Lemma 1 (case *ii*) we have that there is a model M' for D' w.r.t. H' such that $\sum_{w \supseteq \{t'_1,\ldots,t'_m\}} M'(w) = \max\{0, \sum_{i=1}^{m} p(t'_i) - m + 1\}$. We denote this value as p', and consider the case that $p' > 0$ (that is, $p' = \sum_{i=1}^{m} p(t'_i) - m + 1$ as the case that $p' = 0$ can be proved analogously). Since inequality (1) holds for every edge of $HG(D^p, \mathcal{IC})$, the following inequality holds for the tuples of e: $\sum_{i=1..m} p(t'_i) + \sum_{i=1..n} p(t''_i) - m - n + 1 \leq 0$. The quantity $m - \sum_{i=1..m} p(t'_i)$ is equal to $1 - p'$, that is the overall probability assigned by M' to the possible worlds of D' not containing at least one tuple $t'_1, \ldots, t'_m$. Denoting the probability $1 - p'$ as $\overline{p'}$, the above inequality becomes $\sum_{i=1..n} p(t''_i) - n + 1 \leq \overline{p'}$. Owing to Lemma 1 (case *i*), the term on the left-hand side corresponds to $p^{min}(t''_1, \ldots, t''_n)$.

Intuitively enough, this suffices to end the proof, as it means that, if we arrange the tuples $t''_1, \ldots, t''_n$ according to a model M'' for D'' which minimizes the overall probability of the possible worlds of D'' containing $t''_1, \ldots, t''_n$ altogether, the portion of the probability space invested to represent these worlds is less than the portion of the probability space invested by M' to represent the possible worlds of D' not containing at least one tuple among $t'_1, \ldots, t'_m$. For the sake of completeness, we formally show how to obtain a model for D^p w.r.t. $\mathcal{IC}$ starting from M' and M''.

First of all, observe that any interpretation Pr can be represented as a sequence $S(Pr) = (w_1, p1), \ldots, (w_k, p_k)$ where:

- $w_1, \ldots, w_k$ are all the possible worlds such that $Pr(w_i) \neq 0$ for each $i \in [1..k]$;
- $p_1 = Pr(w_1)$;
- for each $i \in [2..n]$ $p_i = p_{i-1} + Pr(w_i)$ (that is, p_i is the cumulative probability of all the possible worlds in $S(M)$ occurring in the positions not greater than i). In particular, this entails that $p_n = 1$.

It is easy to see that many sequences can represent the same interpretation Pr, each corresponding to a different permutation of the set of the possible worlds which are assigned a non-zero probability by Pr.

Consider the model M', and let α be the number of possible worlds which are assigned by M' a non-zero probability and which do not contain at least one tuple among $t'_1, \ldots, t'_m$. Then, take a sequence $S(M')$ such that the first α pairs are possible worlds not containing at least one tuple among $t'_1, \ldots, t'_m$. In this sequence, denoting the generic pair occurring in it as (w'_i, p'_i), it holds that $p'_\alpha = \overline{p'}$.

Analogously, consider the model M'' , and take any sequence $S(M'')$ where the first pair contains the possible world containing all the tuples $t''_1, \ldots, t''_n$. Obviously, denoting the generic pair occurring in $S(M'')$ as (w''_i, p''_i) it holds that $p''_1 = p^{min}(t''_1, \ldots, t''_n)$ is less than or equal to $\overline{p'}$.

Now consider the sequence $S' = (w'''_1, p'''_1), \ldots, ' (w'''_k, p'''_k)$ defined as follows:

- $p'''_1, \ldots, p'''_k$ are the distinct (cumulative) probability values occurring in $S(M')$ and $S(M'')$, ordered by their values;
- for each $i \in [1..k]$, $w'''_i = w'_j \cup w''_l$, where w'_j (resp., w''_l) is the possible world occurring in the left-most pair of $S(M')$ (resp., $S(M'')$) containing a (cumulative) probability value not less than p'''_i.

Consider the function f over the set of possible worlds of D^p defined as follows:

$$f(w) = \begin{cases} 0 & \text{if } w \text{ does not occur in any pair of } S' \\ p'''_1 & \text{if } w \text{ occurs in the first pair of } S' \\ p'''_i - p'''_{i-1} & \text{if } w \text{ occurs in the } i\text{-th pair of } S' (i > 1) \end{cases}$$

It is easy to see that f is an interpretation for D^p. In fact, by construction, it assigns to each possible world of D^p a value in $[0, 1]$, and the sum of the values assigned to the possible worlds is 1. Moreover, the values assigned by f to the possible worlds are compatible with the marginal probabilities of the tuples, since, for each tuple t of D', $\sum_{w'''|t\in w'''} f(w''') = \sum_{w'|t\in w'} M'(w') = p(t)$, as well as for each tuple t of D'', $\sum_{w'''|t\in w'''} f(w''') = \sum_{w''|t\in w''} M''(w'') = p(t)$.

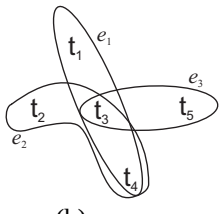
In particular, f is also a model for D^p w.r.t. $\mathcal{IC}$: on the one hand, f assigns 0 to every possible world containing tuples which are conflicting according to H' (this follows from how f was obtained starting from M'). Moreover, f assigns 0 to every possible world containing tuples which are conflicting according to the hyperedge e. In fact, the worlds containing all the tuples $t'_1, \ldots, t'_m, t''_1, \ldots, t''_n$ are assigned 0 by f, since the worlds occurring in S' containing $t''_1, \ldots, t''_n$ do not contain at least one tuple among $t'_1, \ldots, t'_m$ (this trivially follows from the fact that $\overline{p'} > p^{min}(t''_1, \ldots, t''_n)$). The fact that f is a model for D^p w.r.t. $\mathcal{IC}$ means that $D^p \models \mathcal{IC}$. □

The above theorem entails that, if $HG(D^p, \mathcal{IC})$ is a hypertree, then cc can be decided in time $O(|E|\cdot k)$ over $HG(D^p, \mathcal{IC})$, where E is the set of hyperedges of $HG(D^p, \mathcal{IC})$ and k is the maximum arity of the constraints (which bounds the number of nodes in each hyperedge). The number of hyperedges in a hypertree is bounded by the number of nodes $|N|$ (this easily follows from Proposition 1), thus $O(|E| \cdot k) = O(|N| \cdot k)$. Interestingly, even if denial constraints of any arity were allowed, the consistency check could be still accomplished over the conflict hypertree in polynomial time (that is, replacing k with $|N|$, we would get the bound $O(|N|^2)$).

Example 8. *Consider the PDB schme $\mathcal{D}^p$ consisting of relation scheme* Person^p(Name, Age, Parent, Date, City, P) *representing some personal data obtained by integrating various sources. A tuple over* Person^p *refers to a person, and, in particular, attribute* Parent *references the name of one of the parents of the person, while* City *is the city of residence of the person in the date specified in* Date. *Consider the PDB instance D^p consisting of the instance* person^p *of* Person^p *shown in Figure 5(a).*

	Name	*Age*	*Parent*	*Date*	*City*	*P*
t_1	A	40	B	2010	NY	p_1
t_2	A	40	B	2012	Rome	p_2
t_3	A	40	C	2010	NY	p_3
t_4	A	40	D	2010	NY	p_4
t_5	C	30	E	2010	NY	p_5

(a)



(b)

Figure 5. (a) PDB instance D^p; (b) Conflict hypergraph $HG(D^p, \mathcal{IC})$

Assume that $\mathcal{IC}$ consists of the following constraints defined over Person^p*:*

ic_1: $\neg[\, Person(x_1, y_1, z_1, v_1, w_1) \wedge Person(x_1, y_2, z_2, v_2, w_2) \wedge Person(x_1, y_3, z_3, v_3, w_3) \wedge z_1 \neq z_2 \wedge z_1 \neq z_3 \wedge z_2 \neq z_3]$, *imposing that no person has more than 2 parents;*

ic_2: $\neg[\, Person(x_1, y_1, z_1, v_1, w_1) \wedge Person(z_1, y_2, z_2, v_2, w_2) \wedge y_1 > y_2]$, *imposing that no person is older than any of her parents.*

The conflict hypergraph $HG(D^p, \mathcal{IC})$ is shown in Figure 5(b). Here, the conflicting sets e_1, e_2 are originated by violations of ic_1, while e_3 is originated by the violation of ic_2. It is easy to check that $HG(D^p, \mathcal{IC})$ is a hypertree. In particular, observe that set of intersections of e_1 with the other hyper-edges of $HG(D^p, \mathcal{IC})$, that is $\text{Int}(e_1, HG(D^p, \mathcal{IC})) = \{\{t_3\}, \{t_3, t_4\}\}$*, is a matryoshka. Analogously,* $\text{Int}(e_2, HG(D^p, \mathcal{IC}))$ *is matryoshka as well.*

Since $HG(D^p, \mathcal{IC})$ is a hyper-tree, thanks to Theorem 2, we can conclude that D^p is consistent iff the following inequalities hold:

$$p_1 + p_3 + p_4 \leq 2; \qquad p_2 + p_3 + p_4 \leq 2; \qquad p_3 + p_4 \leq 1.$$

□

Note that the condition of Theorem 2 is a necessary condition for consistency in the presence of conflict hypergraphs of any shape, not necessarily hypertrees (in fact, in the proof of the necessary condition of Theorem 2, we did not use the assumption that the conflict hypergraph is a hypertree). The following example shows that this condition is not sufficient in general, in particular when the conflict hypergraph contains "cycles".

Example 9. *Consider the hypergraph $HG(D^p, \mathcal{IC})$ obtained by augmenting the hypertree in Figure 3 with the hyperedge $e_5 = \{t_8, t_9\}$ (whose presence invalidates the acyclicity of the hypergraph). Let the probabilities of $t_1, \ldots, t_9$ be as follows:*

t_i	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9
$p(t_i)$	$\frac{3}{4}$	1	$\frac{3}{4}$	$\frac{3}{4}$	1	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$

Although the condition of Theorem 2 holds for every hyperedge e_i, with i in $[1..5]$*, there is no model of D^p w.r.t. $\mathcal{IC}$. In fact, the overall probability of the possible worlds containing t_8 must be* $1/2$*; due to hyperedges e_4 and e_5, these possible worlds can not contain neither t_6 nor t_9, which must appear together in the remaining possible worlds (since the marginal probability of t_6 and t_9 is equal to the sum of the probabilities of the possible worlds not containing t_8); however, as t_3 can not co-exist with both t_6 and t_9 (due to e_3), it must be in the worlds containing t_8; but, as the overall probability of these worlds is* $1/2$*, they are not sufficient to make the probability of t_3 equal to* $3/4$*.* □

4.1.3. "Cyclic" hypergraphs: cliques and rings

An interesting tractable case which holds even in the presence of cycles in the conflict hypergraph is when the constraints define *buckets* of tuples: buckets are disjoint sets of tuples, such that each pair of tuples in the same bucket are mutually exclusive. The conflict hypergraph describing a set of buckets is simply a graph consisting of disjoint cliques, each one corresponding to a bucket. It is straightforward to see that, in this case, the consistency problem can be decided by just verifying that, for each clique, the sum of the probabilities of the tuples in it is not greater than 1. Observe that the presence of buckets in the conflict hypergraph can be due to key constraints. Thus, what said above implies that cc is tractable in the presence of keys. However, we will be back on the tractability of key constraints in the next section, where we will generalize this tractability result to the presence of one FD per relation.

We now state a more interesting tractability result holding in the presence of "cycles" in the conflict hypergraph.

Theorem 3. *Given an instance D^p of $\mathcal{D}^p$, if $H(D^p, \mathcal{IC}) = \langle N, E \rangle$ is a ring, then $D^p \models \mathcal{IC}$ iff both the following hold: 1) $\forall e \in E,\ \sum_{t \in e} p(t) \leq |e| - 1$; 2) $\sum_{t \in N} p(t) - |N| + \lceil \frac{|E|}{2} \rceil \leq 0$.*

Interestingly, Theorem 3 states that, when deciding the consistency of tuples arranged as a ring in the conflict hypergraph, it is not sufficient to consider the local consistency w.r.t. each hyperedge (as happens in the case of conflict hypertrees), as also a condition involving all the tuples and hyperdges must hold. As an application of this result, consider the case that $H(D^p, \mathcal{IC})$ is the ring whose nodes are t_1, t_2, t_3, t_4 (where: $p(t_1) = p(t_2) = p(t_3) = 1/2$ and $p(t_4) = 1$), and whose edges are: $e_1 = \{t_1, t_2, t_4\}$, $e_2 = \{t_1, t_3, t_4\}$, $e_3 = \{t_2, t_3\}$. It is easy to see that property 1) of Theorem 3 (which is necessary for consistency, as already observed) is satisfied, while property 2) is not (in fact, $\sum_{t \in N} p(t) - |N| + \left\lceil \frac{|E|}{2} \right\rceil = 5/2 - 4 + 2 = 1/2 > 0$), which implies inconsistency. Note that changing $p(t_4)$ to $1/2$ yields consistency.

Remark 2. *Further tractable cases due to the conflict hypergraph.* The tractability results given so far can be straightforwardly merged into a unique more general result: cc is tractable if the conflict hypergraph consists of maximal connected components such that each of them is either a hypertree, a clique, or a ring. In fact, it is easy to see that the consistency can be checked by considering the connected components separately.

4.2. Tractability arising from the syntactic form of the denial constraints

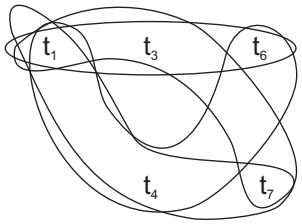
We now address the determination of tractable cases from a different perspective. That is, rather than searching for other properties of the conflict hypergraph guaranteeing that the consistency can be checked in polynomial time, we will search for syntactic properties of denial constraints which can be detected without looking at the conflict hypergraph and which yield the tractability of cc. We start from the following result.

Theorem 4. *If $\mathcal{IC}$ consists of a join-free denial constraint, then* cc *is in PTIME. In particular, $D^p \models \mathcal{IC}$ iff, for each hyperedge e of $HG(D^p, \mathcal{IC})$, it holds that $\sum_{t\in e} p(t) \leq |e| - 1$.*

Example 10. *Consider the PDB scheme consisting of the probabilistic relation scheme* Employee^p(Name, Age, Team, P)*. This scheme is used to represent some (uncertain) personal information about the employees of an enterprise. The uncertain data were obtained starting from anonymized data, and then estimating sensitive information (such as the names of the employees). Assume that the PDB instance D^p obtained this way consists of the instance* employee^p *of* Employee^p *shown in Figure 6(a).*

	Name	*Age*	*Team*	*P*
t_1	P. Jane	35	A	1
t_2	T. Lisbon	25	B	1
t_3	W. Rigsby	40	B	1/2
t_4	K. Cho	40	B	1/2
t_5	G. Van Pelt	22	C	1
t_6	G. Bertram	40	C	1/2
t_7	R. John	40	C	1/2

(a)



(b)

Figure 6. (a) PDB instance D^p; (b) Conflict hypergraph $HG(D^p, \mathcal{IC})$

From some knowledge of the domain, it is known that at least one team among 'A', 'B', 'C' *consists of only young employees, i.e., employees at most* 30*-year old. This corresponds to considering $\mathcal{IC} = \{ic\}$ as the set of denial constraints, where ic is as follows:*

$ic: \neg [\, \text{Employee}(x_1, x_2, \text{'A'}) \wedge \text{Employee}(x_3, x_4, \text{'B'}) \wedge \text{Employee}(x_5, x_6, \text{'C'}) \wedge x_2 > 30 \wedge x_4 > 30 \wedge x_6 > 30 \,]$.

It is easy to see that ic is a join-free denial constraint, thus the consistency of D^p can be decided using Theorem 4. In particular, since $HG(D^p, \mathcal{IC})$ is the hypergraph depicted in Figure 6(b), we have that D^p is consistent if and only if the following inequalities hold:

$$p(t_1) + p(t_3) + p(t_6) \leq 2; \qquad p(t_1) + p(t_3) + p(t_7) \leq 2; \qquad p(t_1) + p(t_4) + p(t_6) \leq 2; \qquad p(t_1) + p(t_4) + p(t_7) \leq 2;$$

As a matter of fact, all these inequalities are satisfied, thus the considered PDB is consistent. In fact, there is a unique model Pr for D^p w.r.t. $\mathcal{IC}$. In particular, Pr assigns probability 1/2 *to each of the possible worlds $w_1 = \{t_1, t_2, t_3, t_4, t_5\}$ and $w_2 = \{t_1, t_2, t_5, t_6, t_7\}$, and probability* 0 *to all the other possible worlds.* □

The result of Theorem 4 strengthens what already observed in the previous section: the arity of constraints is not, *per se*, a source of complexity. In what follows, we show that the arity can become a source of complexity when combined with the presence of join conditions.

Theorem 5. *There is an $\mathcal{IC}$ consisting of a non-join-free denial constraint of arity* 3 *such that* cc *is NP-hard.*

Still, one may be interested in what happens to the complexity of cc for denial constraints containing joins and having arity strictly lower than 3. In particular, since in the proof of Theorem 5 we exploit a ternary EGD to show the *NP*-hardness of cc in the presence of ternary constraints with joins (see Appendix A.3), it is worth investigating what happens when only binary EGDs are considered, which are denial constraints with arity 2 containing joins. The following theorem addresses this case, and states that cc becomes tractable for any $\mathcal{IC}$ consisting of a binary EGD.

Theorem 6. *If $\mathcal{IC}$ consists of a binary EGD, then* cc *is in PTIME.*

Differently from the previous theorems on the tractability of cc, in the statement of Theorem 6, for the sake of presentation, we have not explicitly reported the necessary and sufficient conditions for consistency. In fact, in this setting, deciding on the consistency requires reasoning by cases, and then checking some conditions which are not easy to be defined compactly. However, these conditions can be checked in polynomial time, and the interested reader can find their formal definition in the proof of Theorem 6 (see Appendix A.3).

Binary EGDs can be viewed as a generalization of FDs, involving pairs of tuples possibly belonging to different relations. For instance, over the relation schemes *Student*(*Name*, *Address*, *University*) and *Employee*(*Name*, *Address*, *Firm*), the binary EGD $\neg [Student(x_1, x_2, x_3) \wedge Employee(x_1, x_3, x_4) \wedge x_2 \neq x_3]$ imposes that if a student and an employee are the same person (i.e., they have the same name), then they must have the same address. Thus, an immediate consequence of Theorem 6 is that cc is tractable in the presence of a single FD.

The results presented so far refer to cases where $\mathcal{IC}$ consists of a single denial constraint. We now devote our attention to the case that $\mathcal{IC}$ is not a singleton. In particular, the last tractability result makes the following question arise: *"Is* cc *still tractable when $\mathcal{IC}$ contains several binary EGDs?"*. (Obviously, we do not consider the case of multiple EGDs of any arity, as Theorem 5 states that cc is already hard if $\mathcal{IC}$ merely contains one constraint of this form.) The following theorem provides a negative answer to this question, as it states that cc can be intractable even in the simple case that $\mathcal{IC}$ consists of just two FDs (as recalled above, FDs are special cases of binary EGDs).

Theorem 7. *There is an $\mathcal{IC}$ consisting of* 2 *FDs over the same relation scheme such that* cc *is NP-hard.*

However, the source of complexity in the case of two FDs is that they are defined over the same relation (see the proof of Theorem 7 in Appendix A.3). As a matter of fact, the following theorem states that all the tractability results stated in this section in the presence of only one denial constraint can be extended to the case of multiple denial constraints defined over disjoint sets of relations. Intuitively enough, this derives from the fact that, if the denial constraints involve disjoint sets of relation, the overall consistency can be checked by considering the constraints separately.

Theorem 8. *Let each denial constraint in $\mathcal{IC}$ be join-free or a BEGD. If, for each pair of distinct constraints ic_1,ic_2 in $\mathcal{IC}$, the relation names occurring in ic_1 are distinct from those in ic_2, then* cc *is in PTIME.*

Hence, the above theorem entails that cc is tractable in the interesting case that $\mathcal{IC}$ consists of one FD per relation. In the following theorem, we elaborate more on this case, and specify necessary and sufficient conditions which can be checked to decide the consistency.

Theorem 9. *If $\mathcal{IC}$ consists of one FD per relation, then $HG(D^p, \mathcal{IC})$ is a graph where each connected component is either a singleton or a complete multipartite graph. Moreover, D^p is consistent w.r.t. $\mathcal{IC}$ iff the following property holds: for each connected component C of $HG(D^p, \mathcal{IC})$, denoting the maximal independent sets of C as $S_1, \ldots, S_k$, it is the case that $\sum_{i\in[1..k]} \tilde{p}_i \leq 1$, where $\tilde{p}_i = \max_{t\in S_i} p(t)$.*

We recall that a complete multipartite graph is a graph whose nodes can be partitioned into sets such that an edge exists if and only if it connects two nodes belonging to distinct sets. Each of these sets is a maximal independent set of nodes. For instance, the portion of the graph in Figure 7(b) containing only the nodes t_1, t_2, t_3, t_4, t_5 is a complete multipartite graph whose maximal independent sets are $S_1 = \{t_1, t_2\}$, $S_2 = \{t_3, t_4\}$, $S_3 = \{t_5\}$. The following example shows an application of Theorem 9.

Example 11. *Consider the PDB scheme consisting of the probabilistic relation scheme* Person^p(Name, City, State, P)*, and its instance D^p consisting of the instance* person^p *of* Person^p *shown in Figure 7(a).*

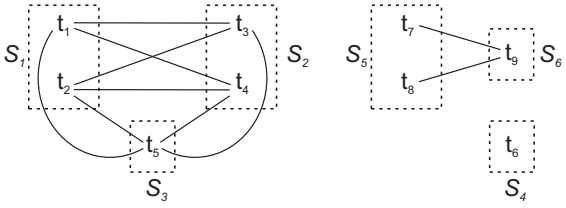
Consider the FD ic: City $\rightarrow$ State*, which can be rewritten as* $\neg[$ Person$(x_1, x_2, x_3) \wedge$ Person$(x_4, x_2, x_5) \wedge x_3 \neq x_5]$. *The conflict hypergraph $HG(D^p, \mathcal{IC})$ is the graph depicted in Figure 7(b). It consists of* 3 *connected components: one of them is a singleton (and corresponds to the maximal independent set S_4), and the other two are the complete multipartite graphs over the maximal independent sets S_1, S_2, S_3 and S_5, S_6, respectively. Theorem 9 says that D^p is consistent if and only if the following three inequalities (one for each connected component of $HG(D^p, \mathcal{IC})$) hold:*

$$\max\{p(t_1), p(t_2)\} + \max\{p(t_3), p(t_4)\} + p(t_5) \leq 1; \qquad p(t_6) \leq 1; \qquad \max\{p(t_7), p(t_8)\} + p(t_9) \leq 1.$$

As a matter of fact, all these inequalities are satisfied, thus the considered PDB is consistent. In fact, there is a model M for D^p w.r.t. $\mathcal{IC}$ assigning probability $1/4$ *to each of the possible worlds $w_1 = \{t_1, t_2, t_6, t_7, t_8\}$, $w_2 = \{t_1, t_6, t_7\}$,*

	Name	*City*	*State*	*P*
t_1	B. Van de Kamp	Sioux City	IA	1/2
t_2	S. Delfino	Sioux City	IA	1/4
t_3	L. Scavo	Sioux City	NE	1/4
t_4	G. Solis	Sioux City	NE	1/4
t_5	E. Britt	Sioux City	SD	1/4
t_6	K. Mayfair	Baltimore	MD	3/4
t_7	R. Perry	Fargo	ND	3/4
t_8	M. A. Young	Fargo	ND	1/4
t_9	K. McCluskey	Fargo	MN	1/4

(a)



(b)

Figure 7. (a) PDB instance D^p; (b) Conflict hypergraph $HG(D^p, \mathcal{IC})$

$w_3 = \{t_3, t_4, t_6, t_7\}$, and $w_4 = \{t_5, t_9\}$, and probability 0 *to all the other possible worlds. The reader can easily check that there are models for D^p w.r.t. $\mathcal{IC}$ other than M.* □

4.3. Tractability implied by conflict-hypergraph properties vs. tractability implied by syntactic forms.

The tractability results stated in sections 4.1 and 4.2 can be viewed as complimentary to each other. In fact, an instance of cc may turn out to be tractable due the syntactic form of the constraints, even if the shape of the conflict hypergraph is none of those ensuring tractability, and vice versa. For instance, in the case that $\mathcal{IC}$ consists of a join-free denial constraint or a binary EGD, it is easy to see that the conflict hypergraph may not be a hypertree or a ring, but cc is nevertheless tractable due to theorems 2 and 3. Vice versa, if $\mathcal{IC}$ contains two FDs per relation or a ternary denial constraints with joins (which, potentially, are hard cases, due to theorems 5 and 7), cc may turn out to be tractable, if the way the data combine with the constraints yields a conflict hypergraph which is a hypertree or a ring (see theorems 2 and 3).

On the whole, the tractability results presented in sections 4.1 and 4.2 can be used conjunctively when addressing cc: for instance, one can start by examining the constraints and check whether they conform to a tractable syntactic form, and, if this is not the case, one can look at the conflict hypergraph and check whether its structure entails tractability.

5. Querying PDBs under constraints

As explained in the previous section, given a PDB D^p in the presence of a set $\mathcal{IC}$ of integrity constraints, not all the interpretations of D^p are necessarily models w.r.t. $\mathcal{IC}$. If D^p is consistent w.r.t. $\mathcal{IC}$, there may be exactly one model (Case 2 of the motivating example), or more (Case 3 of the same example). In the latter case, given that all the models satisfy all the constraints in $\mathcal{IC}$, there is no reason to assume one model more reasonable than the others (at least in the absence of other knowledge not encoded in the constraints). Hence, when querying D^p, it is "cautious" to answer to queries by taking into account all the possible models for D^p w.r.t. $\mathcal{IC}$. In this section, we follow this argument and introduce a *cautious querying paradigm* for conjunctive queries, where query answers consist of tuples associated with probability ranges: given a query Q posed over D^p, the range associated with a tuple t in the answer of Q contains every probability with which t would be returned as an answer of Q if Q were evaluated separately on every model of D^p. In what follows, we first introduce the formal definition of conjunctive query in the probabilistic setting, and introduce its semantics according to the above-discussed cautious paradigm. Then, we provide our contributions on the characterization of the problem of computing query answers.

A (conjunctive) query over a PDB schema $\mathcal{D}^p$ is written as a (conjunctive) query over its deterministic part $det(\mathcal{D}^p)$. Thus, it is an expression of the form:
$Q(\vec{x}) = \exists \vec{z}.\ R_1(\vec{y}_1) \wedge \cdots \wedge R_m(\vec{y}_m) \wedge \phi(\vec{y}_1, \ldots, \vec{y}_m)$, where:

– $R_1, \ldots, R_m$ are name of relations in $det(\mathcal{D}^p)$;
– $\vec{x}$ and $\vec{z}$ are tuples of variables, having no variable common;

– $\vec{y}_1,\ldots,\vec{y}_m$ are tuples of variables and constants such that every variable in any $\vec{y}_i$ occurs in either $\vec{x}$ or $\vec{z}$, and vice versa;
– $\phi(\vec{y}_1,\ldots,\vec{y}_m)$ is a conjunction of built-in predicates, each of the form $\alpha \diamond \beta$, where α and β are either variables in $\vec{y}_1,\ldots,\vec{y}_m$ or constants, and $\diamond \in \{=,\neq,\leq,\geq,<,>\}$.

A query Q will be said to be *projection-free* if $\vec{z}$ is empty.

The semantics of a query Q over a PDB D^p in the presence of a set of integrity constraints $\mathcal{IC}$ is given in two steps. First, we define the answer of Q w.r.t. a single model M of D^p. Then, we define the answer of Q w.r.t. D^p, which summarizes all the answers of Q obtained by separately evaluating Q over every model of D^p . Obviously, we rely on the assumption that D^p is consistent w.r.t. $\mathcal{IC}$, thus $\mathcal{M}(D^p,\mathcal{IC})$ is not empty.

The answer of Q over a model M of D^p w.r.t. $\mathcal{IC}$ is the set $Ans^M(Q,D^p,\mathcal{IC})$ of pairs of the form $\langle\vec{t}, p_Q^M(\vec{t})\rangle$ such that:

– $\vec{t}$ is a ground tuple such that $\exists w \in pwd(D^p)$ *s.t.* $w \models Q(\vec{t})$;
– $p_Q^M(\vec{t}) = \sum_{w\in pwd(D^p)\wedge w\models Q(\vec{t})} M(w)$ is the overall probability of the possible worlds where $Q(\vec{t})$ evaluates to true,

where $w \models Q(\vec{t})$ denotes that $Q(\vec{t})$ evaluates to true in w.

In general, there may be several models for D^p, and the same tuple $\vec{t}$ may have different probabilities in the answers evaluated over different models. Thus, the overall answer of Q over D^p is defined in what follows as a summarization of all the answers of Q over all the models of D^p.

Definition 4 (Query answer)**.** *Let Q be a query over $\mathcal{D}^p$, and D^p an instance of $\mathcal{D}^p$. The answer of Q over D^p is the set* $\mathrm{Ans}(Q,D^p,\mathcal{IC})$ *of pairs $\langle\vec{t},[p^{min},p^{max}]\rangle$, where:*

– $\exists M \in \mathcal{M}(D^p,\mathcal{IC})$ *s.t. $\vec{t}$ is a tuple in* $\mathrm{Ans}^M(Q,D^p,\mathcal{IC})$*;*
– $p^{min} = \min_{M\in\mathcal{M}(D^p,\mathcal{IC})}\left\{p_Q^M(\vec{t})\right\}, \qquad p^{max} = \max_{M\in\mathcal{M}(D^p,\mathcal{IC})}\left\{p_Q^M(\vec{t})\right\}.$

Hence, each tuple $\vec{t}$ in $Ans(Q,D^p,\mathcal{IC})$ is associated with an interval $[p^{\min},p^{\max}]$, whose extremes are, respectively, the minimum and maximum probability of $\vec{t}$ in the answers of Q over the models of D^p. Examples of answers of a query are reported in the motivating example. In the following, we say that $\vec{t}$ is an answer of Q with minimum and maximum probabilities $p^{\min}$ and $p^{\max}$ if $\langle\vec{t},[p^{\min},p^{\max}]\rangle \in Ans(Q,D^p,\mathcal{IC})$.

The following proposition gives an insight on the semantics of query answers, as it better explains the meaning of the probability range associated with each tuple occurring in the set of answers of a query. That is, it states that, taken any pair $\langle\vec{t},[p^{\min},p^{\max}]\rangle$ in $Ans(Q,D^p,\mathcal{IC})$, every value p inside the interval $[p^{\min},p^{\max}]$ is "meaningful", in the sense that there is at least one model for which $\vec{t}$ is an answer of Q with probability p. Considering this property along the fact that the boundaries $p^{\min}$, $p^{\max}$ are the minimum and maximum probabilities of $\vec{t}$ as an answer of Q (which follows from Definition 4), we have that $[p^{\min},p^{\max}]$ is the tightest interval containing all the probabilities of $\vec{t}$ as an answer of Q, and is dense (every value inside it corresponds to a probability of $\vec{t}$ as an answer of Q).

Proposition 2. *Let Q be a query over $\mathcal{D}^p$, and D^p an instance of $\mathcal{D}^p$. For each pair $\langle\vec{t},[p^{\min},p^{\max}]\rangle$ in* $\mathrm{Ans}(Q,D^p,\mathcal{IC})$*, and each probability value $p \in [p^{\min},p^{\max}]$, there is a model M of D^p w.r.t. $\mathcal{IC}$ such that $\langle\vec{t},p\rangle \in \mathrm{Ans}^M(Q,D^p,\mathcal{IC})$.*

Proof. We first introduce a system $S(\mathcal{D}^p,\mathcal{IC},D^p)$ of linear (in)equalities whose solutions one-to-one correspond to the models of D^p w.r.t. $\mathcal{IC}$. For every $w_i \in pwd(D^P)$, let v_i be a variable ranging over the domain of rational numbers. The variable v_i will be used to represent the probability assigned to w_i by an interpretation of D^p. The system of linear (in)equalities $S(\mathcal{D}^p,\mathcal{IC},D^p)$ is as follows:

$$\begin{cases} \forall t \in D^p,\ \sum_{i|w_i\in pwd(D^p)\wedge t\in w_i} v_i = p(t) & (e1)\\ \sum_{i|w_i\in pwd(D^p)\wedge w_i\not\models\mathcal{IC}} v_i = 0 & (e2)\\ \sum_{i|w_i\in pwd(D^p)} v_i = 1 & (e3)\\ \forall w_i \in pwd(D^p),\ v_i \geq 0 & (e4)\end{cases}$$

The first $|D^p|$ equalities $(e1)$ in $S(\mathcal{D}^p,\mathcal{IC},D^p)$ encode the fact that, for each tuple t in the PDB instance, the sum of the probabilities assigned to the worlds containing the tuple t must be equal to the marginal probability of t. The

subsequent two equalities ($e2$), ($e3$), along with the inequalities ($e4$) imposing that the probabilities v_i assigned to each possible world are non-negative, entail that the probability assigned to any world violating $\mathcal{IC}$ is 0, as well as that the probabilities assigned to all the possible worlds sum up to 1.

It is easy to see that every solution s of $S(\mathcal{D}^p, \mathcal{IC}, D^p)$ one-to-one corresponds to a model Pr for D^p w.r.t. $\mathcal{IC}$, where $Pr(w_i)$ is equal to $v_i[s]$, i.e., the value of v_i in s.

We now consider the system of linear (in)equalities $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ obtained by augmenting the set of (in)equalities in $S(\mathcal{D}^p, \mathcal{IC}, D^p)$ with the following equality:

$$v^* = \sum_{i|w_i \in pwd(D^p) \wedge w_i \models \vec{t}} v_i$$

where v^* is a new variable symbol not appearing in $S(\mathcal{D}^p, \mathcal{IC}, D^p)$.

Obviously, every solution s of $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ still one-to-one corresponds to a model Pr for D^p w.r.t. $\mathcal{IC}$ such that, for each possible world $w_i \in pwd(D^p)$, $Pr(w_i)$ is equal to $v_i[s]$, and $v^*[s]$ (the value of v^* in s) is equal to the sum of the probabilities assigned by Pr to the possible worlds where $\vec{t}$ is an answer of Q. Therefore, $p^{\min}$ (resp. $p^{\max}$) is the solution of the following linear programming problem $LP(S^*)$:

$$\begin{array}{l} \text{minimize (resp. maximize) } v^* \\ \text{subject to } S^*(\mathcal{D}^p, \mathcal{IC}, D^p) \end{array}$$

Since the feasible region shared by the min- and max- variants of $LP(S^*)$ is defined by linear inequalities only, it follows that it is a convex polyhedron. Hence, the following well-known result [42] can be exploited: "*given two linear programming problem LP_1 and LP_2 minimizing and maximizing the same objective function f over the same convex feasible region S, respectively, it is the case that for any value v belonging to the interval $[v^{\min}, v^{\max}]$, whose extreme values are the optimal solutions of LP_1 and LP_2, respectively, there is a solution s of S such that v is the value taken by f when evaluated over s*". This result entails that, for every probability value $p \in [p^{\min}, p^{\max}]$ taken by the objective function v^* of $LP(S^*)$, there is a feasible solution s of $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ such that $p = v^*[s]$. Hence, the statement follows from the fact that every solution of $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ one-to-one corresponds to a model for D^p w.r.t. $\mathcal{IC}$. □

The definition of query answers with associated ranges is reminiscent of the treatment of aggregate queries in inconsistent databases [4]. In that framework, the consistent answer of an aggregate query *Agg* is a range $[v_1, v_2]$, whose boundaries represent the minimum and maximum answer which would be obtained by evaluating *Agg* on at least one repair of the database. However, the consistent answer is not, in general, a dense interval: for instance, it can happen that there are only two repairs, one corresponding to v_1 and one to v_2, while the values between v_1 and v_2 can not be obtained as answers on any repair.

In the rest of this section, we address the evaluation of queries from two standpoints: we first consider a decision version of the query answering problem, and then we investigate the query evaluation as a search problem. In the following, besides assuming that a database schema $\mathcal{D}^p$ and a set of constraints $\mathcal{IC}$ of fixed size are given, we also assume that queries over $\mathcal{D}^p$ are of fixed size. Thus, all the complexity results refer to data complexity.

5.1. Querying as a decision problem

In the classical "deterministic" relational setting, the decision version of the query answering problem is commonly defined as the *membership problem* of deciding whether a given tuple belongs to the answer of a given query. In our scenario, tuples belong to query answers with some probability range, thus it is natural to extend this definition to our probabilistic setting in the following way.

Definition 5 (Membership Problem (MP))**.** *Given a query Q over $\mathcal{D}^p$, an instance D^p of $\mathcal{D}^p$, a ground tuple $\vec{t}$, and the constants k_1 and k_2 (with $0 \leq k_1 \leq k_2 \leq 1$), the membership problem is deciding whether $\vec{t}$ is an answer of Q with minimum and maximum probabilities p^{min} and p^{max} such that $p^{min} \geq k_1$ and $p^{max} \leq k_2$.*

Hence, solving MP can be used to decide whether a given tuple is an answer with a probability which is at least k_1 and not greater than k_2. Observe that Definition 5 collapses to the classical definition of membership problem when

data are deterministic: in fact, asking whether a tuple belongs to the answer of a query posed over a deterministic database corresponds to solving MP over the same database with $k_1 = k_2 = 1$.

From the results in [35], where an entailment problem more general than MP was shown to be in co*NP* (see Section 7), it can be easily derived that MP is in co*NP* as well. The next theorems (which are preceded by a preliminary lemma) determine two cases when this upper bound on the complexity is tight.

Lemma 2. *Let Q be a conjunctive query over $\mathcal{D}^p$, D^p an instance of $\mathcal{D}^p$, and $\vec{t}$ an answer of Q having minimum probability $p^{\min}$ and maximum probability $p^{\max}$. Let m be the number of tuples in D^p plus 3 and a be the maximum among the numerators and denominators of the probabilities of the tuples in D^p. Then $p^{\min}$ and $p^{\max}$ are expressible as fractions of the form $\frac{\eta}{\delta}$, with $0 \leq \eta \leq (ma)^m$ and $0 < \delta \leq (ma)^m$.*

Theorem 10 (Lower bound of MP)**.** *There is at least one conjunctive query containing projection for which MP is coNP-hard, even if $\mathcal{IC}$ is empty.*

Proof. We show a LOGSPACE reduction from the consistency checking problem (cc) in the presence of binary denial constraints, which is NP-hard (see Theorem 7), to the complement of the membership problem ($\overline{\text{MP}}$).

Let $\langle \mathcal{D}^p_{\text{cc}}, \mathcal{IC}_{\text{cc}}, D^p_{\text{cc}} \rangle$ be an instance of cc. We construct an equivalent instance $\langle \mathcal{D}^p_{\overline{\text{MP}}}, \mathcal{IC}_{\overline{\text{MP}}}, D^p_{\overline{\text{MP}}}, Q, t_\emptyset, k_1, k_2 \rangle$ of $\overline{\text{MP}}$ as follows.

- $\mathcal{D}^p_{\overline{\text{MP}}}$ consists of relation schemas $R^p(tid, P)$ and $S^p(tid_1, tid_2, P)$;
- $\mathcal{IC}_{\overline{\text{MP}}} = \emptyset$, that is, no constraint is assumed on $\mathcal{D}^p_{\overline{\text{MP}}}$;
- $D^p_{\overline{\text{MP}}}$ is the instance of $\mathcal{D}^p_{\overline{\text{MP}}}$ which contains, for each tuple $t \in D^p_{\text{cc}}$, the tuple $R^p(id(t), p(t))$, where $id(t)$ is a unique identifier associated to the tuple t. Moreover, $D^p_{\overline{\text{MP}}}$ contains, for each pair of tuples t_1, t_2 in D^p_{cc} which are conflicting w.r.t. $\mathcal{IC}_{\text{cc}}$, the tuple $S^p(id(t_1), id(t_2), 1)$.
- $Q = \exists x, y\, R(x) \wedge R(y) \wedge S(x, y)$;
- $t_\emptyset$ is the empty tuple;
- the lower bound k_1 of the minimum probability of $t_\emptyset$ as answer of Q is set equal to $k_1 = \frac{1}{(ma)^m}$, where m is the number of tuples in $D^p_{\overline{\text{MP}}}$ plus 3, and a the maximum among the numerators and denominators of the probabilities of the tuples in $D^p_{\overline{\text{MP}}}$;
- the upper bound k_2 of the maximum probability of $t_\emptyset$ as answer of Q is set equal to 1.

Obviously, the $\overline{\text{MP}}$ instance returns true iff the minimum probability that $t_\emptyset$ is an answer to Q over $D^p_{\overline{\text{MP}}}$ is (strictly) less than k_1.

It is easy to see that every interpretation of D^p_{cc} (the database in the cc instance) corresponds to a unique interpretation of $D^p_{\overline{\text{MP}}}$ (the database in the $\overline{\text{MP}}$ instance), and vice versa. Observe that $D^p_{\overline{\text{MP}}}$ is consistent, since the set of constraints considered in the $\overline{\text{MP}}$ instance is empty.

We show now that the above-considered cc and $\overline{\text{MP}}$ instances are equivalent, that is, the cc instance is true iff the $\overline{\text{MP}}$ instance is true. On the one hand, if the cc instance is true, then there is at least is one model Pr_{cc} for D^p_{cc} w.r.t. $\mathcal{IC}_{\text{cc}}$ (that is, Pr_{cc} assigns probability 0 to every possible world w which contains tuples which are conflicting according to $\mathcal{IC}_{\text{cc}}$). It is easy to see that evaluating Q on the corresponding interpretation $Pr_{\overline{\text{MP}}}$ of $\overline{\text{MP}}$ yields probability 0 for the empty tuple $t_\emptyset$. Hence, the $\overline{\text{MP}}$ instance is true in this case.

On the other hand, if the $\overline{\text{MP}}$ instance is true, then the minimum probability that $t_\emptyset$ is an answer of Q must be less than $\frac{1}{(ma)^m}$. Since $\frac{1}{(ma)^m}$ is the smallest non-zero value that can be assumed by the minimum probability of $t_\emptyset$ (see Lemma 2), this implies that the minimum probability that $t_\emptyset$ is an answer of Q is 0. This means that there is a model $Pr_{\overline{\text{MP}}}$ that assigns probability 0 to every possible world w which contains three tuples $R(x_1)$, $R(y_1)$ and $S(x_2, y_2)$ with $x_1 = x_2$ and $y_1 = y_2$. It is easy to see that the corresponding interpretation Pr_{cc} is a model for D^p_{cc} w.r.t. $\mathcal{IC}_{\text{cc}}$, as it assigns probability 0 to every possible world which contains conflicting tuples. Hence the cc instance is true in this case. □

The above theorem establishes that the type of the query, and in particular that fact that it contains projection, is an important source of complexity making MP hard, irrespectively of the constraints considered. For projection-free queries, the next theorem states that MP remains hard even if only binary constraints are considered.

Theorem 11 (Lower bound of MP)**.** *There is at least one projection-free conjunctive query and a set $\mathcal{IC}$ consisting of only binary constraints for which* MP *is coNP-hard.*

We recall that, when addressing MP, we assume that the database is consistent w.r.t. the constraints. Thus, the hardness results for MP do not derive from any source of complexity inherited by MP from CC. On the whole, theorems 10 and 11 suggest that MP has at least two sources of complexity: the type of query (the fact that the query contains projection or not), and the form of the constraints.

Once some sources of complexity of MP have been identified, the problem is worth addressing of determining tractable cases. We defer this issue after the characterization of the query evaluation as a search problem, since, as it will be clearer in what follows, the conditions yielding tractability of the latter problem also ensure the tractability of MP.

5.2. Querying as a search problem

Viewed a search problem, the query answering problem (QA) is the problem of computing the set $Ans(Q, D^p, \mathcal{IC})$. The complexity of this problem is characterized as follows.

Theorem 12. QA *is in FP^{NP} and is $FP^{NP[\log n]}$-hard.*

The fact that QA is in FP^{NP} means that our "cautious" query evaluation paradigm is not more complex than the query evaluation based on the independence assumption, which has been shown in [11] to be complete for $\#P$ (which strictly contains FP^{NP}, assuming $P \neq NP$). On the other hand, the hardness for $FP^{NP[\log n]}$ is interesting also because it tightens the characterization given in [35] of the more general entailment problem for probabilistic logic programs containing a general form of probabilistic rules (*conditional rules*). Specifically, in [35], the above-mentioned entailment problem was shown to be in FP^{NP}, but no lower bound on its data complexity was stated. Thus, our result enriches the characterization in [35], as it implies that $FP^{NP[\log n]}$ is a lower bound for the entailment problem for probabilistic logic programs under data complexity even in the presence of rules much simpler than conditional rules. More details are given in Section 7, where we provide a more thorough comparison with [35]. However, finding the tightest characterization for QA remains an open problem, as it might be the case that QA is complete for either $FP^{NP[\log n]}$ or FP^{NP}. We conjecture that none of these cases holds (thus a characterization of QA tighter than ours can not be provided), thus QA is likely to be in the "limbo" containing the problems in FP^{NP} but not in $FP^{NP[\log n]}$, without being hard for the former (this limbo is non-empty if $P \neq NP$ [30]).

5.3. Tractability results

In this section, we show some sufficient conditions for the tractability of the query evaluation problem, which hold for both its decision and search versions. When stating our results, we refer to QA only, as its tractability implies that of MP (as MP is straightforwardly reducible to QA).

Again, we address the tractability from two standpoints: we will show sufficient conditions which regard either *a*) the shape of the conflict hypergraph, or *b*) the syntactic form of the constraints. Specifically, we focus on finding islands of tractability when queries are projection-free and either the conflict hypergraph collapses to a graph – as for direction *a*), or the constraints are binary – as for direction *b*). These are interesting contexts, since Theorem 11 entails that MP (and, thus, also QA) is, in general, hard in these cases (indeed, Theorem 11 implicitly shows the hardness for the case of conflict hypergraphs collapsing to graphs, as, in the presence of binary constraints, the conflict hypergraph is a graph).

The next result goes into direction *a*), as it states that, for projection-free queries, QA is tractable if the conflict hypergraph is a graph satisfying some structural properties.

Theorem 13. *For projection-free conjunctive queries,* QA *is in* PTIME *if $HG(D^p, \mathcal{IC})$ is a graph where each maximal connected component is either a tree or a clique.*

The polynomiality result stated above is rather straightforward in the case that each connected component is a clique, but is far from being straightforward in the presence of connected components which are trees. Basically, when the conflict hypergraph is a tree, the tractability derives from the fact that, for any conjunction of tuples, its minimum (or, equivalently, maximum) probability can be evaluated as the solution of an instance of a linear programming problem. In particular, differently from the "general" system of inequalities used in the proof of Proposition 2 (where

the variables corresponds to the possible worlds, thus their number is exponential in the number of tuples), here we can define a system of inequalities where both the number of inequalities and variables depend only on the arity of the query (which is constant, as we address data complexity). We do not provide an example of the form of this system of inequalities, as explaining the correctness of the approach on a specific case is not easier than proving its validity in the general case. Thus, the interested reader is referred to the proof of Theorem 13 reported in Appendix A.6 for more details.

The following result goes into direction of locating tractability scenarios arising from the syntactic form of the constraints, as it states that, if $\mathcal{IC}$ consists of one FD for each relation scheme, the evaluation of projection-free queries is tractable.

Theorem 14. *For projection-free conjunctive queries,* QA *is in* PTIME *if $\mathcal{IC}$ consists of at most one FD per relation scheme.*

Proof. We consider the case that $\mathcal{IC}$ contains one relation and one FD only, as the general case (more relations, and one FD per relation) follows straightforwardly. Let the denial constraint *ic* in $\mathcal{IC}$ be the following FD over relation scheme R: $X \rightarrow Y$, where X, Y are disjoint sets of attributes of R. We denote as r the instance of R in the instance of QA. Constraint *ic* implies a partition of r into disjoint relations, each corresponding to a different combination of the values of the attributes in X in the tuples of r. Taken one of this combinations $\vec{x}$ (i.e., $\vec{x} \in \Pi_X(r)$), we denote the corresponding set of tuples in this partition as $r(\vec{x})$. That is, $r(\vec{x}) = \{t \in r | \Pi_X(t) = \vec{x}\}$. In turn, for each $r(\vec{x})$, *ic* partitions it into disjoint relations, each corresponding to a different combinations of the values of the attributes in Y. Taken one of this combinations $\vec{y}$ (i.e., $\vec{y} \in \Pi_Y(r(\vec{x}))$), we denote the corresponding set of tuples in this partition as $r(\vec{x}, \vec{y})$.

Given this, constraint *ic* entails that the conflict hypergraph is a graph with the following structure: there is an edge (t_1, t_2) iff $\exists \vec{x}, \vec{y_1}, \vec{y_2}$, with $\vec{y_1} \neq \vec{y_2}$, such that $t_1 \in r(\vec{x}, \vec{y_1})$ and $t_2 \in r(\vec{x}, \vec{y_2})$.

Now, consider any conjunction of tuples $T = t_1, \ldots, t_n$. The probability of T as an answer of the query q specified in the instance of QA can be computed as follows. First, we partition $\{t_1, \ldots, t_n\}$ according to the maximal connected components of the conflict hypergraph. This way we obtain the disjoint subsets $T_1, \ldots, T_k$ of $\{t_1, \ldots, t_n\}$, where each T_i corresponds to a maximal connected component of the conflict hypergraph, and contains all the tuples of $\{t_1, \ldots, t_n\}$ which are in this component. The minimum and maximum probabilities of T as answer of q can be obtained by computing the minimum and maximum probability of each set T_i, and then combining them using the well known Frechet-Hoeffding formulas (reported also in the appendix as Fact 2), which give the minimum and maximum probabilities of a conjunction of events among which no correlation is known (in fact, since $T_1, \ldots, T_k$ correspond to distinct connected components, they can be viewed as pairwise uncorrelated events).

Then, it remains to show how the minimum and maximum probabilities of a single T_i can be computed. We consider the case that T_i contains at least two tuples (otherwise, the minimum and the maximum probabilities of T_i coincide with the marginal probability of the unique tuple in T_i). If $\exists t_\alpha, t_\beta \in T_i$ $\exists \vec{x}, \vec{y_1}, \vec{y_2}$ such that $t_\alpha \neq t_\beta$ and $\vec{y_1} \neq \vec{y_2}$ and $t_\alpha \in r(\vec{x}, \vec{y_1})$, while $t_\beta \in r(\vec{x}, \vec{y_2})$, then the minimum and maximum probabilities of T_i are both 0 (since $\{t_\alpha, t_\beta\}$ is a conflicting set). Otherwise, it is the case that all the tuples in T_i share all the values $\vec{x}$ for the attributes X, and the same values $\vec{y}$ for the attributes Y. Due to the structure of the conflict hypergraph, it is easy to see that this implies that the tuples in T_i can be distributed in any way in the portion of the probability space which is not invested to represent the tuples having the same values $\vec{x}$ for X, but combinations for Y other than $\vec{y}$. The size of this probability space is $S = 1 - \sum_{\vec{y^*} \neq \vec{y}} \max\{p(t) | t \in r(\vec{x}, \vec{y^*})\}$. Hence, the minimum and maximum probabilities of T_i are:

$$p^{\min} = \max\left\{0, \sum_{t \in T_i} p(t) - |T_i| + S\right\}; \qquad p^{\max} = \min\{p(t) \mid t \in T_i\}.$$

The first formula is an easy generalization of the corresponding formula for the minimum probability given in Lemma 1 to the case of a probability space of a generic size less than 1. The second formula derives from the above-recalled Frechet-Hoeffding formulas, and from the fact that the database is consistent (we recall that we rely on this assumption when addressing the query evaluation problem). □

Again, observe that the last two results are somehow complementary: it is easy to see that there are FDs yielding conflict hypergraphs not satisfying the sufficient condition of Theorem 13, as well as conflict hypergraphs which are trees generated by some "more general" denial constraint, not expressible as a set of FDs over distinct relations.

6. Extensions of our framework

Some extensions of our framework are discussed in what follows. In particular, for each extension, we show its impact on our characterization of the fundamental problems addressed in the paper.

6.1. Tuples with uncertain probabilities

All the results stated in this paper can be trivially extended to the case that tuples are associated with ranges of probabilities, rather than single probabilities (as happens in several probabilistic data models, such as [31, 35]).

Obviously, all the hardness results for CC, MP, QA hold also for this variant, since considering tuples with single probabilities is a special case of allowing tuples associated with range of probabilities.

As regards CC, both the membership in *NP* and the extendability of the tractable cases straightforwardly derive from the fact that, as only denial constraints are considered, deciding on the consistency of an assignment of ranges of probabilities can be accomplished by looking only at the minimum probabilities of each range.

As regards MP and QA, the fact that the complexity upper-bounds do not change follows from the results in [35]. Finally, it can be shown, with minor changes to the proof of Theorem 13, that MP and QA are still tractable under the hypotheses on the shape of the conflict hypergraph stated in this theorem. We refer the interested reader to Appendix A.7, where a hint is given on how the proof of Theorem 13 can be extended to deal with tuples with uncertain probabilities. The extension of the tractability results for MP and QA regarding the syntactic forms of the constraints is even simpler, and can be easily understood after reading the proofs of these results.

6.2. Associating constraints with probabilities.

Another interesting extension consists in allowing constraints to be assigned probabilities. In our vision, constraints should encode some certain knowledge on the data domain, thus they should be interpreted as deterministic. However, this extension can be interesting at least from a theoretical point of view, or when constraints are derived from some elaboration on historical data [18]. Thus, the point becomes that of giving a semantics to the probability assigned to the constraints. The semantics which seems to be the most intuitive is as follows: "*A constraint with probability p forbidding the co-existence of some tuples is satisfied if there is an interpretation where the overall probability of the possible worlds satisfying the constraint is at least p*". This means that the condition imposed by the constraint must hold in a portion of size p of the probability space, while nothing is imposed on the remaining portion of the probability space.

Starting from this, we first discuss the impact of associating constraints with probabilities on our results about CC. First of all, it is easy to see that there is a reduction from any instance *Prob-cc* of the variant of CC with probabilistic constraints to an equivalent instance *Std-cc* of the standard version of CC. Basically, this reduction constructs the conflict hypergraph H(*Std-cc*) of *Std-cc* as follows: denoting the conflict hypergraph of *Prob-cc* as H(*Prob-cc*), each hyperedge $e \in H$(*Prob-cc*) (with probability $p(e)$) is transformed into a hyperedge e' of H(*Std-cc*) which consists of the same nodes in e plus a new node with probability $p(e)$. On the one hand, the existence of this reduction suffices to state that also the probabilistic version of CC is *NP*-complete. On the other hand, it is worth noting that applying this reduction yields a conflict hypergraph H(*Std-cc*) with the same "shape" as H(*Prob-cc*), except that each hyperedge has one new node, belonging to no other hyperedge: hence, if H(*Prob-cc*) is a hypertree (resp., a ring), then H(*Std-cc*) is a hypertree (resp., a ring) too. This means that all the tractability results given for CC concerning the shapes of the conflict hypergraph hold also when stated directly on its probabilistic version. However, this does not suffice to extend the tractability results for CC regarding the syntactic forms of the constraints, as in the considered cases the conflict hypergraph may not be a hypertree or a ring. Thus, the extension of the tractability results on the syntactic forms is deferred to future work.

As regards MP and QA, the arguments used in the discussion of the previous extension can be used to show that our lower and upper bounds still hold for the variants of these problems allowing probabilistic constraints. As for the tractability results, in Appendix A.7, a more detailed discussion is provided explaining how the proof of Theorem 13 (which deal with conflict hypergraphs where each maximal connected componenent is either a clique or a tree) can be extended to deal with probabilistic constraints. The extension of the tractability result for FDs stated in Theorem 14 is deferred to future work.

6.3. Assuming pairs of tuples as independent unless this contradicts the constraints

As observed in the introduction, in some cases, rejecting the assumption of independence for some groups of tuples may be somehow "overcautious". For instance, if we consider further tuples pertaining to a different hotel in the introductory example (where constraints involve tuples over the same hotel), it may be reasonable to assume that these tuples encode events independent from those pertaining hotel 1.

A naive way of extending our framework in this direction is that of assuming every pair of tuples which are not explicitly "correlated" by some constraint as independent from one another. This means considering as independent any two tuples t_1, t_2 such that there is no hyperedge in the conflict hypergraph containing both of them. However, this strategy can lead to wrong interpretations of the data. For instance, consider the case of Example 3, where each of the three tuples t_1, t_2, t_3 has probability $1/2$, and two (ground) constraints are defined over them: one forbidding the co-existence of t_1 with t_2, and the other forbidding the co-existence of t_2 with t_3. As observed in Example 3, the combination of these two constraints implicitly enforces the co-existence of t_1 with t_3. Hence, the fact that t_1 and t_3 are not involved in the same (ground) constraint does not imply that these two tuples can be considered as independent from one another.

However, it is easy to see that if two tuples are not connected through any path in the conflict hypergraph, assuming independence among them does not contradict the constraints in any way. Hence, a cautious way of incorporating the independence assumption in our framework is the following: any two tuples are independent from one another iff they belong to distinct maximal connected components of the conflict hypergraph.

If this model is adopted, nothing changes in our characterization of the consistency checking problem. In fact, it is easy to see that an instance of cc is equivalent to an instance of the variant of cc where independence is assumed among maximal connected components of the conflict hypergraph. This trivially follows from the fact that, if a PDB D^p is consistent according to the original framework, all the possible interpretations combining the models of the maximal connected components are themselves models of D^p, and the set of these interpretations contains also the interpretation corresponding to assuming independence among the maximal connected components.

As regards the query evaluation problem, adopting this variant of the framework makes QA #P-hard (as QA becomes more general than the problem of evaluating queries under the independence assumption [11]). However, all our tractability results for projection-free queries still hold. In fact, the probability of $t_1, \ldots, t_n$ as an answer of a query can be obtained as follows. First, the set $T = \{t_1, \ldots, t_n\}$ is partitioned into the (non-empty) sets $S_1, \ldots, S_k$ which correspond to distinct maximal connected components of the conflict hypergraph, and where each S_i consists of all the tuples in T belonging to the connected component corresponding to S_i. Then, the minimum and maximum probabilities of each S_i are computed (in *PTIME*, when our sufficient conditions for tractability hold), by considering each S_i separately. Finally, the independence assumption among the tuples belonging to distinct maximal components is exploited, so that the minimum (resp., maximum) probability of $t_1, \ldots, t_n$ is evaluated as the product of the so obtained minimum (resp., maximum) probabilities of $S_1, \ldots, S_k$.

7. Related work

We separately discuss the related work in the AI and DB literature.

AI setting. The works in the AI literature related to ours are mainly those dealing with probabilistic logic. The problem of integrating probabilities into logic was first addressed (though pretty informally) in [39]. Then, in [22] the PSAT problem was formalized as the satisfiability problem in a propositional fragment of the logic discussed in [39], and shown to be *NP*-complete. In [17], a more general probabilistic propositional logic than that in [22] was defined, which enables algebraic relations to be specified among the probabilities of propositional formulas (such as "*the probability of* $\phi_1 \wedge \phi_2$ *is twice that of* $\phi_3 \vee \phi_4$). [17] mainly focuses on the satisfiability problem, showing that it is *NP*-complete (thus generalizing the result on PSAT of [22]). However, it provides no tractability result (whose investigation is our main contribution in the study of the corresponding consistency problem). Up to our knowledge, most of the works devising techniques for efficiently solving the satisfiability problem (such as [27, 34]) rely on translating it into a Linear Programming instance and using some heuristics, which do not guarantee polynomial-bounded complexity. Thus, the only works determining provable polynomial cases of probabilistic satisfiability are [2, 22]. As for [22], we refer the reader to the discussions in Section 4 (right after Definition 3) and at the begininning of Section 4.1. As regards [2], it is related to our work in that it showed that PSAT is tractable if the hypergraph of the formula (which

corresponds to our conflict hypergraph) is a hypertree. However, the notion of hypertree in [2] is very restrictive, as it relies on a notion of acyclicity much less general than the γ-acyclicity used here. In fact, even the simple hypergraph consisting of $e_1 = \{t_1, t_2, t_3\}$, $e_2 = \{t_2, t_3, t_4\}$ is not viewed in [2] as a hypertree, since it contains at least one cycle, such as $t_1, e_1, t_2, e_2, t_3, e_1, t_1$ (note that, in our framework, this would not be a cycle). Basically, hypertrees in [2] are special cases of our hypertrees, as they require distinct hyperedges to have at most one node in common. Hence, our result strongly generalizes the forms of conflict hypergraphs over which cc turns out to be tractable according to the result of [2] on PSAT.

The entailment problem (which corresponds to our query answering problem) was studied both in the propositional [34] and in the (probabilistic-)logic-programming setting [35, 38, 37]. The relationship between these works and ours is in the fact that they deal with knowledge bases where rules and facts can be associated with probabilities. Intuitively, imposing constraints over a PDB might be simulated by a probabilistic logic program, where tuples are encoded by (probabilistic) facts and constraints by (probabilistic) rules with probability 1. However, not all the above-cited probabilistic-logic-programming frameworks can be used to simulate our framework: for instance, [38, 37] use rules which can not express our constraints. On the contrary, the framework in [35] enables pretty general rules to be specified, that is *conditional rules* of the form $(H|B)[p_1, p_2]$, where H and B are classical open formulas, stating that the probability of the formula $H \wedge B$ is between p_1 and p_2 times the probability of B. Obviously, any denial constraint *ic* can be written as a conditional rule of the form $(H|true)[1, 1]$, where H is the open formula in *ic*. In the presence of conditional rules, [35] characterizes the complexity of the satisfiability and the entailment problems. The novelty of our contribution w.r.t. that of [35] derives from the specific database-oriented setting considered in our work. In particular, as regards the consistency problem, our tractable cases are definitely a new contribution, as [35] does not determine polynomially-solvable instances. As regards the query answering problem, our contribution is relevant from several standpoints. First, we provide a lower bound of the membership problem by assuming that the database is consistent: this is a strong difference with [35], where the decisional version of the entailment problem has been addressed without assuming the satisfiability of the knowledge base, thus the satisfiability checking is used as a source of complexity when deciding the entailment. Second, we have characterized the lower bound of the membership problem w.r.t. two specific aspects, which make sense in a database-perspective and were not considered in [35]: the presence of projection in the query (Theorem 10) and the type of denial constraints (Theorem 11). Third, [35] did not prove any lower bound for the data complexity of the search version of the entailment problem. Indeed, it provided an FP^{NP}-hardness result only under combined complexity (assuming all the knowledge base as part of the input, while we consider constraints of fixed size) and exploiting the strong expressiveness of conditional rules, which enable also constraints not expressible by denial constraints to be specified. Hence, in brief, our Theorem 12 shows that constraints simpler than conditional constraints suffice to get an $FP^{NP[\log n]}$-hardness of the entailment for probabilistic logic programs, even under data complexity. Finally, our tractable cases of the query evaluation problem, up to our knowledge, are not subsumed by any result in the literature, and depict islands of tractability also for the more general entailment problem studied in [35].

DB setting. The database research literature contains several works addressing various aspects related to probabilistic data, and a number of models have been proposed for their representation and querying. In this section, we first summarize the most important results on probabilistic databases relying on the independence assumption (which, obviously, is somehow in contrast with allowing integrity constraints to be specified over the data, thus making these works marginally related to ours). Then, we focus our attention on other works, which are more related to ours as they allow some forms of correlations among data to be taken into account when representing and querying data.

As regards the works relying on the independence assumption, the problem of efficiently evaluating (conjunctive) queries was first studied in [11], where it was shown that this problem is #P-hard in the general case of queries without self-joins, but can be solved in polynomial time for queries admitting a particular evaluation plan (namely, *safe plan*). Basically, a safe plan is obtained by suitably pushing the projection in the query expression, in order to extend the validity of the independence assumption also to the partial results of the query. The results of [11] were extended in [10, 14, 13, 24, 41]. Specifically, in [14], a technique was presented for computing safe plans on disjoint-independent databases (where only tuples belonging to different buckets are considered as independent). In [13] and [10], the dichotomy theorem of [11] was extended to deal with conjunctive queries with self-joins and unions of conjunctive queries, respectively. In [41], it was shown that a polynomial-time evaluation can be accomplished also with query plans with any join ordering (not only those orderings required by safe plans). Finally, in [24], a

technique was presented enabling the determination of efficient query plans even for queries admitting no safe plan (this is allowed by looking at the database instance to decide the most suitable query plan, rather than looking only at the database schema).

The problem of dealing with probabilistic data when correlations are not known (and independence may not be assumed) was addressed in [31]. Here, an algebra for querying probabilistic data was introduced, as well as a system called *ProbView*, which supports the evaluation of algebraic expressions by returning answers associated with probability intervals. However, the query evaluation is based on an extensional semantics and no integrity constraints encoding domain knowledge were considered.

One of the first works investigating a suitable model for representing correlations among probabilistic data is [23], where *probabilistic c-tables* were introduced. In this framework, whose rationale is also at the basis of the PDB *MayBMS* [28], correlations are expressed by associating tuples with boolean formulas on random variables, whose probability functions are represented in a table. However, in this approach, only one interpretation for the database is considered (the one deriving from assuming the random variables independent from one another), and it is not suitable for simulating the presence of integrity constraints on the data when the marginal probabilities of the tuples are known. Similar differences, such as that of assuming only one interpretation, hold between our framework and that at the basis of *Trio* [5, 1], where incomplete and probabilistic data are modeled by combining the possibility of specifying buckets of tuples with the association of each tuple with its lineage (expressed as the set of tuples from which each tuple derived). In particular, in [1] an extension of Trio is proposed which aims at better managing the epistemic uncertainty (i.e., the information about uncertainty is itself incomplete). Here, the semantics of *generalized uncertain databases* is given in terms of a Dempster-Shafer mass distribution over the powerset of the possible worlds (this collapses to the case of a PDB with one probability distribution, if the mass distribution is defined over every single possible world). Further approaches to representing rich correlations and querying the data are those in [43, 32, 26], where correlations among data are represented according to some graphical models (such as PGMs, junction trees, AND/XOR trees). In these approaches, correlations are detected while data are generated and, in some sense, they are data themselves: the database consists of a graph representing correlations among events, so that the marginal distributions of tuples are not explicitly represented, but derive from the correlations encoded in the graph. This is a strong difference with our framework, where a PDB is a set of tuples associated with their marginal probabilities, and constraints can be imposed by domain experts with no need of taking part to the data-acquisition process. Moreover, in [43, 32, 26], independence is assumed between tuples for which a correlation is not represented in the graph of correlations. On the contrary, our query evaluation model relies on a "cautious" paradigm, where no assumption is made between tuples not explicitly correlated by the constraints. In [12], the problem of evaluating queries over probabilistic views under integrity constraints (functional and inclusion dependencies) and in the presence of statistics on the cardinality of the source relations was considered. In this setting, when evaluating query answers and their probabilities, all the possible values of the attribute values of the original relations must be taken into account, and this backs the use of the Open World Assumption (as the original relations may contain attribute values which do not occur in the views). Under this assumption, queries are evaluated over the interpretation of the data having the maximum entropy among all the possible models.

All the above-cited works assume that the correlations represented among the data are consistent. In [29], the problem was addressed of querying a PDB when integrity constraints are considered a posteriori, thus some possible worlds having non-zero probability under the independence assumption may turn out to be inconsistent. In this scenario, queries are still evaluated on the unique interpretation entailed by the independence assumption, but the possible worlds are assigned the probabilities *conditioned* to the fact that what entailed by the constraint is true. That is, in the presence of a constraint Γ, the probability $P(Q)$ of a query Q is evaluated as $P(Q|\Gamma)$, which is the probability of Q assuming that Γ holds. This corresponds to evaluating queries by augmenting them with the constraints, thus it is a different way of interpreting the constraints and queries from the semantics adopted in our paper, where constraints are applied on the database. The same spirit as this approach is at the basis of [9], where specific forms of integrity constraints in the special case of probabilistic XML data are taken into account by considering a single interpretation, conditioned on the constraints.

8. Conclusions and Future work

We have addressed two fundamental problems dealing with PDBs in the presence of denial constraints: the *consistency checking* and the *query evaluation* problem. We have thoroughly studied the complexity of these problems, characterizing the general cases and pointing out several tractable cases.

There exist a number of interesting directions for future work. First of all, the cautious querying paradigm will be extended to deal with further forms of constraints. This will allow for enriching the types of correlations which can be expressed among the data, and this may narrow the probability ranges associated with the answers (in fact, for queries involving tuples which are not involved in any denial constraint, the obtained probability ranges may be pretty large, and of limited interest for data analysis).

Another interesting direction for future work is the identification of other tractable cases of the consistency checking and the query evaluation problems. As regards the consistency checking problem, we conjecture that polynomial-time strategies can be devised when the conflict hypergraph exhibits a limited degree of cyclicity (as a matter of fact, we have shown that this problem is feasible in linear time not only for hypertrees, but also for rings, which have limited cyclicity as well). A possible starting point is investigating the connection between the consistency checking problem (viewed as evaluating the (dual) lineage of the constraint query - see Remark 1) and the model checking problem of Boolean formulas. The connection between lineage evaluation and model checking has been well established mainly for the cases of tuple-independent PDBs [40, 25]. In fact, in this setting, it has been shown that, as it happens for checking Boolean formulas, the probability of a lineage can be evaluated by compiling it into a *Binary Decision Diagram* - BDD [36], and then suitably processing the diagram. Specifically, if the lineage (or, equivalently, the Boolean formula to be checked) L can be compiled into a particular case of BDDs (such as *Read-Once* or *Ordered* BDD), the lineage evaluation (as well as the formula verification) can be accomplished as the result of a traversal of the BDD, in time linear w.r.t. the diagram size. Hence, in all the cases where L can be compiled into a *Read-Once* or an *Ordered* BDD of polynomial size, L can be evaluated in polynomial time. One of the most general result about the compilability of Boolean formulas into *Ordered* BDDs was stated in [19], where it was shown that any CNF expression over n variables whose hypergraph of clauses has bounded treewidth ($< k$) admits an equivalent ordered BDD of size $O(n^{k+1})$. Then, the point becomes devising a mechanism for exploiting an Ordered BDD equivalent to a Boolean formula f to evaluate the probability of f, when neither independence nor precise correlations can be assumed among the terms of f. Up to our knowledge, this topic has not been investigated yet, and we plan to address it in future work. If it turned out that, under no assumption on the way terms are correlated, the probability of formulas can be evaluated by traversing their equivalent Ordered BDDs, then the above-cited result of [19] would imply other tractable cases of our consistency checking problem. However, our results on hypertrees and rings would be still of definite interest, as we have found that in these cases the consistency checking problem can be solved in linear time, while the construction of the ordered BDD is $O(n^{k+1})$. Moreover, our results show that the consistency checking problem over hypertrees and rings is still polynomially solvable (actually, in quadratic time) in the case that the cardinality of hyperedges is not known to be bounded by constants (see the discussion right after Theorem 2), which does not always correspond to structures having bounded treewidth.

Finally, our framework can be exploited to address the problem of repairing data and extracting reliable information from inconsistent PDBs. This research direction is somehow related to [3], where the evaluation of clean answers over *deterministic* databases which are inconsistent due to the presence of duplicates is accomplished by encoding the inconsistent database into a PDB adopting the bucket independent model. Basically, in this PDB, probabilities are assigned to tuples representing variants of the same tuple, and these variants are grouped in buckets. However, the so obtained PDB is consistent, thus this approach is not a repairing framework for inconsistent PDBs, but is a technique for getting clean answers over inconsistent deterministic databases after rewriting queries into "equivalent" queries over the corresponding consistent PDBs. A more general repairing problem in the probabilistic setting has been recently addressed in [33], where a strategy based on deleting tuples has been proposed, "inspired" by the common approaches for inconsistent deterministic databases [6]. We envision a different repairing paradigm, which addresses a source of inconsistency which is typical of the probabilistic setting: inconsistencies may arise from wrong assignments to the marginal probabilities of tuples, due to limitations of the model adopted for encoding uncertain data into probabilistic tuples. In this perspective, a repairing strategy based on properly updating the probabilities of the tuples (possibly by adapting frameworks for data repairing in the deterministic setting based on attribute updates [20, 21, 45]) seems to be the most suitable choice.

Acknowledgements. *We are grateful to the anonymous reviewers of an earlier conference submission of a previous version of this paper for their fruitful suggestions (one especially for pointing out the reduction of* cc *to PSAT), as well as Thomas Lukasiewicz, for insightful discussions about his work [35], and Francesco Scarcello, for valuable comments about our work.*

References

[1] P. Agrawal, J. Widom, Generalized uncertain databases: First steps, in: Proc. 4th Int. VLDB workshop on Management of Uncertain Data (MUD), pp. 99–111.
[2] K.A. Andersen, D. Pretolani, Easy cases of probabilistic satisfiability, Annals of Mathematics and Artificial Intelligence (AMAI) 33 (2001).
[3] P. Andritsos, A. Fuxman, R.J. Miller, Clean answers over dirty databases: A probabilistic approach, in: Proc. 22nd Int. Conf. on Data Engineering (ICDE), p. 30.
[4] M. Arenas, L.E. Bertossi, J. Chomicki, X. He, V. Raghavan, J. Spinrad, Scalar aggregation in inconsistent databases, Theoretical Computer Science 296 (2003) 405–434.
[5] O. Benjelloun, A.D. Sarma, A.Y. Halevy, J. Widom, Uldbs: Databases with uncertainty and lineage, in: Proc. 32nd Int. Conf. on Very Large Data Bases (VLDB), pp. 953–964.
[6] L. Bertossi, Database Repairing and Consistent Query Answering, Morgan & Claypool Publishers, 2011.
[7] G. Boole, An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities, Macmillan, London, 1854.
[8] J. Chomicki, J. Marcinkowski, S. Staworko, Computing consistent query answers using conflict hypergraphs, in: Proc. 2004 Int. Conf. on Information and Knowledge Management (CIKM), pp. 417–426.
[9] S. Cohen, B. Kimelfeld, Y. Sagiv, Incorporating constraints in probabilistic XML, ACM Transactions on Database Systems 34 (2009).
[10] N.N. Dalvi, K. Schnaitter, D. Suciu, Computing query probability with incidence algebras, in: Proc. 29th Symp. on Principles of Database Systems (PODS), pp. 203–214.
[11] N.N. Dalvi, D. Suciu, Efficient query evaluation on probabilistic databases, in: Proc. 30th Int. Conf. on Very Large Data Bases (VLDB), pp. 864–875.
[12] N.N. Dalvi, D. Suciu, Answering queries from statistics and probabilistic views, in: Proc. 31st Int. Conf. on Very Large Data Bases (VLDB), pp. 805–816.
[13] N.N. Dalvi, D. Suciu, The dichotomy of conjunctive queries on probabilistic structures, in: Proc. 26th Symp. on Principles of Database Systems (PODS), pp. 293–302.
[14] N.N. Dalvi, D. Suciu, Management of probabilistic data: foundations and challenges, in: Proc. 26th Symp. on Principles of Database Systems (PODS), pp. 1–12.
[15] A. D'Atri, M. Moscarini, On the recognition and design of acyclic databases, in: Proc. 3rd Symp. on Principles of Database Systems (PODS), pp. 1–8.
[16] R. Fagin, Degrees of acyclicity for hypergraphs and relational database schemes, Journal of the ACM 30 (1983).
[17] R. Fagin, J. Halpern, N. Megiddo, A logic for reasoning about probabilities, Information and Computation (IC) 87 (1990) 78–128.
[18] F. Fassetti, B. Fazzinga, Fox: Inference of approximate functional dependencies from xml data, in: 2nd Int. DEXA Workshop on XML Data Management Tools and Techniques (XANTEC), pp. 10–14.
[19] A. Ferrara, G. Pan, M.Y. Vardi, Treewidth in verification: Local vs. global, in: Proc. 12th Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR), pp. 489–503.
[20] S. Flesca, F. Furfaro, F. Parisi, Preferred database repairs under aggregate constraints, in: Proc. 1st Int. Conf. on Scalable Uncertainty Management (SUM), pp. 215–229.
[21] S. Flesca, F. Furfaro, F. Parisi, Querying and repairing inconsistent numerical databases, ACM Transactions on Database Systems 35 (2010).
[22] G.F. Georgakopoulos, D.J. Kavvadias, C.H. Papadimitriou, Probabilistic satisfiability, Journal of Complexity 4 (1988) 1–11.
[23] T.J. Green, V. Tannen, Models for incomplete and probabilistic information, in: Proc. 2006 EDBT Workshop on Inconsistency and Incompleteness in Databases (IIDB), pp. 278–296.
[24] A.K. Jha, D. Olteanu, D. Suciu, Bridging the gap between intensional and extensional query evaluation in probabilistic databases, in: Proc. 13th Int. Conf. on Extending Database Technology (EDBT), pp. 323–334.
[25] A.K. Jha, D. Suciu, Knowledge compilation meets database theory: compiling queries to decision diagrams, in: Proc. 14th Int. Conf. on Database Theory (ICDT), pp. 162–173.
[26] B. Kanagal, A. Deshpande, Lineage processing over correlated probabilistic databases, in: Proc. 2010 Int. Conf. on Management of Data (SIGMOD), pp. 675–686.
[27] D.J. Kavvadias, C.H. Papadimitriou, A linear programming approach to reasoning about probabilities, Annals of Mathematics and Artificial Intelligence 1 (1990) 189–205.
[28] C. Koch, MayBMS: A system for managing large uncertain and probabilistic databases, Managing and Mining Uncertain Data Ch. 9 (2009).
[29] C. Koch, D. Olteanu, Conditioning probabilistic databases, Proceedings of the VLDB Endowment (PVLDB) 1 (2008) 313–325.
[30] M.W. Krentel, The complexity of optimization problems, Journal of Computer and System Sciences 36 (1988) 490–509.
[31] L.V.S. Lakshmanan, N. Leone, R.B. Ross, V.S. Subrahmanian, Probview: A flexible probabilistic database system, ACM Transactions on Database Systems 22 (1997) 419–469.
[32] J. Li, A. Deshpande, Consensus answers for queries over probabilistic databases, in: Proc. 28th Symp. on Principles of Database Systems (PODS), pp. 259–268.

[33] X. Lian, L. Chen, S. Song, Consistent query answers in inconsistent probabilistic databases, in: Proc. 2010 Int. Conf. on Management of Data (SIGMOD), pp. 303–314.
[34] T. Lukasiewicz, Probabilistic deduction with conditional constraints over basic events, Journal of Artificial Intelligence Research (JAIR) 10 (1999) 199–241.
[35] T. Lukasiewicz, Probabilistic logic programming with conditional constraints, ACM Transactions on Computational Logic 2 (2001) 289–339.
[36] C. Meinel, T. Theobald, Algorithms and Data Structures in VLSI Design, Springer-Verlag, 1998.
[37] R.T. Ng, Semantics, consistency, and query processing of empirical deductive databases, IEEE Transactions on Knowledge and Data Engineering (TKDE) 9 (1997) 32–49.
[38] R.T. Ng, V.S. Subrahmanian, Probabilistic logic programming, Information and Computation 101 (1992) 150–201.
[39] N.J. Nilsson, Probabilistic logic, Artificial Intelligence 28 (1986) 71–87.
[40] D. Olteanu, J. Huang, Using obdds for efficient query evaluation on probabilistic databases, in: Proc. 2nd Int. Conf. on Scalable Uncertainty Management (SUM), pp. 326–340.
[41] D. Olteanu, J. Huang, C. Koch, SPROUT: Lazy vs. eager query plans for tuple-independent probabilistic databases, in: Proc. 25th Int. Conf. on Data Engineering (ICDE), pp. 640–651.
[42] C.H. Papadimitriou, K. Steiglitz, Combinatorial Optimization: Algorithms and Complexity, Dover Publications, Inc, Mineola, New York, 1998.
[43] P. Sen, A. Deshpande, Representing and querying correlated tuples in probabilistic databases, in: Proc. 23rd Int. Conf. on Data Engineering (ICDE), pp. 596–605.
[44] P. Sen, A. Deshpande, L. Getoor, Prdb: managing and exploiting rich correlations in probabilistic databases, VLDB Journal 18 (2009) 1065–1090.
[45] J. Wijsen, Database repairing using updates, ACM Transactions on Database Systems 30 (2005) 722–768.

Appendix A. Proofs

In this appendix we report the proofs of the theorems whose statement have been provided and commented in the main body of the paper. Furthermore, the appendix contains some new lemmas which are exploited in these proofs.

Appendix A.1. Proofs of Theorem 1, Proposition 1, and Lemma 1

Theorem 1. (Complexity of cc) cc *is NP-complete*.

Proof. The membership of cc in *NP* has been already proved in the core of the paper, where a reduction from cc to *PSAT* has been described. As regards the hardness, it follows from Theorem 7 (or, equivalently, from Theorem 5), whose proof is given in Section Appendix A.3. □

We now report a property of γ-acyclic hypergraphs from [15], which will be used in the proof of Proposition 1.

Fact 1. [15] *Let $H = \langle N, E\rangle$ be a hypertree. There exists at least one hyperedge $e \in E$ such that at least one of the following conditions hold:*

1. *$e \cap \mathrm{N}(H^{-\{e\}})$ is a set of edge equivalent nodes;*
2. *there exists $e' \in E$ such that $e' \neq e$ and $e \cap \mathrm{N}(H^{-\{e,e'\}}) = e' \cap \mathrm{N}(H^{-\{e,e'\}})$.*

Moreover, $H^{-\{e\}}$ is still a hypertree.

Proposition 1. *Let $H = \langle N, E\rangle$ be a hypertree. Then, there is at least one hyperedge $e \in E$ such that $Int(e, H)$ is a matryoshka. Moreover, $H^{-\{e\}}$ is still a hypertree.*

Proof. Reasoning by induction on the number of hyperedges in E, we prove that there is a total ordering $e_1, \cdots, e_n$ of the edges in E such that all the following conditions hold for each $i \in [1..n]$:

1. either $e_i \cap N(H^{-\{e_1,\cdots,e_{i-1}\}})$ is a set of edge equivalent nodes, or there exists $e' \in E(H^{-\{e_1,\cdots,e_{i-1}\}})$ such that $e' \neq e$ and $e \cap N(H^{-\{e,e'\}}) = e' \cap N(H^{-\{e,e'\}})$;
2. $H^{-\{e_1,\cdots,e_i\}}$ is a hypertree;
3. $Int(e_i, H^{-\{e1,\cdots,e_{i-1}\}})$ is a matryoshka.

The base case ($|E| = 1$) is straightforward. In order to prove the induction step, we reason as follows. Since H is a hypertree, Fact 1 implies that there is a node e such that 1) either $e \cap N(H^{-\{e\}})$ is a set of edge equivalent nodes, or there exists $e' \in E$ such that $e' \neq e$ and $e \cap N(H^{-\{e,e'\}}) = e' \cap N(H^{-\{e,e'\}})$, and 2) $H^{-\{e\}}$ is a hypertree.

From the inductive hypothesis, since $H^{-\{e\}}$ is a hypertree, there exists a total ordering $e_1, \cdots, e_{n-1}$ of the nodes in $E - \{e\}$ such that for each $i \in [1..n-1]$ conditions 1, 2 and 3 are satisfied w.r.t. $H^{-\{e\}}$.

If $Int(e, H)$ is a matryoshka, then the total ordering $e, e_1, \cdots, e_{n-1}$ of the nodes in E satisfies conditions 1, 2 and 3 for every edge in the sequence thus the statement is proved in this case.

Otherwise, since $Int(e, H)$ is not a matryoshka then $e \cap N(H^{-\{e\}})$ is not a set of edge equivalent nodes. Hence, since e satisfies the conditions of Fact 1 then there exists $e_j \in \{e_1, \cdots, e_{n-1}\}$ such that $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$.

We now consider separately the following two cases:

Case 1): there is $k \in [1..j-1]$ such that $e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}})$.

Case 2): there is no $k \in [1..j-1]$ such that $e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}})$.

We first prove Case 1). Let $k \in [1..j-1]$ be the smallest index such that $e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}})$. We consider the total ordering of the edges of E obtained by inserting e immediately before e_k in $e_1, \cdots, e_{n-1}$, i.e., $e_1, \cdots, e_{k-1}, e, e_k, \cdots, e_j, \cdots, e_{n-1}$.

We first prove that for each $i \in [1..k-1]$ conditions 1, 2 and 3 still hold. For each $i \in [1..k-1]$ one of the following cases occur:

- $e_i \cap e_j = \emptyset$. In this case since $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$, it is straightforward to see that conditions 1, 2 and 3 hold.

- $e_i \cap e_j \neq \emptyset$ and $e_i \cap N(H^{-\{e,e_1,\cdots,e_{i-1}\}})$ is a set of edge equivalent nodes. Since $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$, $e_i \cap e_j \neq \emptyset$ and e_j is an edge of $H^{-\{e,e_1,\cdots,e_{i-1}\}}$ then $e_i \cap N(H^{-\{e,e_1,\cdots,e_{i-1}\}}) = e_i \cap N(H^{-\{e_1,\cdots,e_{i-1}\}})$. Therefore, the nodes in $e_i \cap N(H^{-\{e_1,\cdots,e_{i-1}\}})$ are edge equivalent w.r.t $H^{-\{e_1,\cdots,e_{i-1}\}}$ too. Hence, conditions 1, 2 and 3 hold.

- $e_i \cap e_j \neq \emptyset$ and there is an $h \in [i+1..n-1]$, with $h \neq j$, such that $e_i \cap N(H^{-\{e,e_1,\cdots,e_i,e_h\}}) = e_h \cap N(H^{-\{e,e_1,\cdots,e_i,e_h\}})$. Since, e_j is and edge of $H^{-\{e,e_1,\cdots,e_i,e_h\}}$ and $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$ it holds that $e_i \cap N(H^{-\{e,e_1,\cdots,e_i,e_h\}}) = e_i \cap N(H^{-\{e_1,\cdots,e_i,e_h\}}) = e_h \cap N(H^{-\{e_1,\cdots,e_i,e_h\}})$. Hence conditions 1, 2 and 3 hold in this case too.

Observe that, in the last two cases mentioned above the fact that $Int(e_i, H^{-\{e1,\cdots,e_{i-1}\}})$ is a matryoshka follows from the fact that $e_i \cap N(H^{-\{e1,\cdots,e_{i-1}\}}) = e_i \cap N(H^{-\{e,e1,\cdots,e_{i-1}\}})$ and $e_i \cap e = e_i \cap e_j$. Moreover, conditions 1, 2 and 3 still hold for each $i \in [k..n-1]$ since they are not changed w.r.t. the inductive hypothesis.

As regards the edge e, it is easy to see that conditions 1 and 2 are satisfied since e_j appears after e in the total ordering $e_1, \cdots, e_{k-1}, e, e_k, \cdots, e_j, \cdots, e_{n-1}$.

We now prove that condition 3 holds for e. We know from the induction hypothesis that $Int(e_k, H^{\{e,e_1,\cdots,e_{k-1}\}})$ is a matryoshka. However, since $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$ and $j > k$ then $Int(e_k, H^{\{e,e_1,\cdots,e_{k-1}\}}) = Int(e_k, H^{\{e_1,\cdots,e_{k-1}\}})$. Since, $e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}})$ and $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$ it holds that

$$e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = e \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}).$$

Therefore the set of nodes in $e \cap N(H^{-\{e_1,\cdots,e_{k-1}\}})$ can be partitioned in three sets N, N', N'' such that:

- $N = e_k \cap N(H^{-\{e,e_1,\cdots,e_k,e_j\}}) = \bigcup_{S \in Int(e_k, H^{\{e,e_1,\cdots,e_{k-1}\})}} S$,
- $N' = e_k \cap e_j - N$, and
- $N'' = e \cap e_j - N' - N$.

Hence, it is easy to see that $Int(e, H^{\{e_1,\cdots,e_{k-1}\}}) = Int(e_k, H^{\{e,e_1,\cdots,e_{k-1}\}}) \cup \{N \cup N' \cup \{N \cup N' \cup N''\}$. Therefore, $Int(e, H^{\{e_1,\cdots,e_{k-1}\}})$ is a matryoshka. Hence, the proof for Case 1) is completed.

We now prove Case 2). We consider the total ordering of the edges of E obtained by inserting e immediately before e_j in $e_1, \cdots, e_{n-1}$, i.e., $e_1, \cdots, e_{j-1}, e, e_j, \cdots, e_{n-1}$. It is easy to see that we can prove that for each $i \in [1..j-1]$ conditions 1, 2 and 3 are satisfied applying the same reasoning applied in Case 1) in order to prove that for each $i \in [1..k-1]$ conditions 1, 2 and 3 hold. Analogously to the proof of Case 1) it is straightforward to see that conditions 1, 2 and 3 still hold for each $i \in [j..n-1]$ since they are not changed w.r.t. the inductive hypothesis.

As regards the edge e, it is easy to see that conditions 1 and 2 are satisfied since e_j appears after e in the total ordering $e_1, \cdots, e_{j-1}, e, e_j, \cdots, e_{n-1}$.

To complete the proof we show that condition 3 holds for e in this case. From the induction hypothesis, we know that it is the case that $Int(e_j, H^{\{e,e_1,\cdots,e_{j-1}\}})$ is a matryoshka. However, since $e \cap N(H^{-\{e,e_j\}}) = e_j \cap N(H^{-\{e,e_j\}})$ then $e \cap N(H^{-\{e,e_1,\cdots,e_j\}}) = e_j \cap N(H^{-\{e,e_1,\cdots,e_j\}})$, and it holds that the set of nodes in $e \cap N(H^{-\{e_1,\cdots,e_{k-1}\}})$ can be partitioned in the sets N and N' such that:

- $N = e_j \cap N(H^{-\{e,e_1,\cdots,e_j\}}) = \bigcup_{S \in Int(e_j, H^{\{e,e_1,\cdots,e_{k-1}\})}} S$,
- $N' = e_k \cap e_j - N$.

It is easy to see that the following holds $Int(e, H^{\{e_1,\cdots,e_{j-1}\}}) = Int(e_j, H^{\{e,e_1,\cdots,e_{k-1}\}}) \cup \{N \cup N'\}$. Therefore, $Int(e, H^{\{e_1,\cdots,e_{j-1}\}})$ is a matryoshka, which completes the proof for Case 2) and the proof of the proposition. □

Before providing the proof of Lemma 1, we report a well-known result on the minimum and maximum probability of the conjunction of events among which no correlation is known, taken from [7].

Fact 2. *Let E_1, E_2 be a pair of events such that their marginal probabilities $p(E_1)$, $p(E_2)$ are known, while no correlation among them is known. Then, the minimum and maximum probabilities of the event $E_1 \wedge E_2$ are as follows: $p^{min}(E_1 \wedge E_2) = \max\{0, p(E_1) + p(E_2) - 1\}$; and $p^{max}(E_1 \wedge E_2) = \min\{p(E_1), p(E_2)\}$.*

The formulas reported above are also known as Frechet-Hoeffding formulas. In Lemma 1, we generalize the formula for the minimum probability, and adapt it to our database setting.

Lemma 1. *Let D^p be an instance of $\mathcal{D}^p$ consistent w.r.t. $\mathcal{IC}$, T a set of tuples of D^p, and $H = HG(D^p, \mathcal{IC})$. If either i) the tuples in T are pairwise disconnected in H, or ii) $Int(T, H)$ is a matryoshka, then $p^{min}(T) = \max\{0, \sum_{t\in T} p(t) - |T| + 1\}$. Otherwise, this formula provides a lower bound for $p^{min}(T)$.*

Proof. Case *i*): In the case that $t_1, \ldots, t_n$ are pairwise disconnected in the conflict hypergraph, the formula for $p^{min}(t_1, \ldots, t_n)$ can be proved by induction on n, considering as base case the formula for the minimum probability of a pair of events reported in Fact 2.

Case *ii*): We prove an equivalent formulation of the statement over the same instance of D^p: *"Let T be a set of nodes of $H = HG(D^p, \mathcal{IC})$ such that $Int(T, H)$ is a matryoshka. Let $T^n = t_1, \ldots, t_n$ be a sequence consisting of the nodes of T ordered as follows: $i > j \implies s(t_i) \supseteq s(t_j)$, where $s(t_i)$ is the maximal set in $Int(T, H)$ containing t_i. Then, $p^{min}(t_1, \ldots, t_n) = \max\left\{0, \sum_{i=1}^{n} p(t_i) - n + 1\right\}$"*. That is, we consider the nodes in T suitably ordered, as this will help us to reason inductively.

We reason by induction on the length of the sequence T^n. The base case ($n = 1$) trivially holds, as, for any tuple t, $p^{min}(t) = p(t)$. We now prove the induction step: we assume that the property holds for any sequence of the considered form of length $n - 1$, and prove that this implies that the property holds for sequences of n nodes.

From induction hypothesis, we have that the property holds for the subsequence $T^{n-1} = t_1, \ldots, t_{n-1}$ of T^n. That is, there is a model M for D^p w.r.t. $\mathcal{IC}$ such that $\sum_{w\supseteq\{t_1,\ldots,t_{n-1}\}} M(w) = \max\{0, \sum_{i=1}^{n-1} p(t_i) - (n-1) + 1\}$. We show how, starting from M, a model M' can be constructed such that $\sum_{w\supseteq\{t_1,\ldots,t_{n-1},t_n\}} = \max\{0, \sum_{i=1}^{n} p(t_i) - n + 1\}$, which is the formula reported in the statement for $p^{min}(t_1, \ldots, t_n)$. According to M, the set of possible worlds of D^p can be partitioned into:

- $W(t_1 \wedge \cdots \wedge t_{n-1} \wedge t_n)$: the set of possible worlds containing all the tuples $t_1, \ldots, t_{n-1}, t_n$;
- $W(\neg(t_1 \wedge \cdots \wedge t_{n-1}) \wedge t_n)$: the set of possible worlds containing t_n, but not containing at least one among $t_1, \ldots, t_{n-1}$;
- $W(t_1 \wedge \cdots \wedge t_{n-1}, \neg t_n)$: the set of possible worlds containing all the tuples $t_1, \ldots, t_{n-1}$, but not containing t_n;
- $W(\neg(t_1 \wedge \cdots \wedge t_{n-1}) \wedge \neg t_n)$: the set of possible worlds not containing t_n and not containing at least one tuple among $t_1, \ldots, t_{n-1}$.

For the sake of brevity, the set of worlds defined above will be denoted as W, W', W'', W''', respectively. In the following, given a set of possible worlds $\mathcal{W}$, we denote as $M(\mathcal{W})$ the overall probability assigned by M to the worlds in $\mathcal{W}$, i.e., $M(\mathcal{W}) = \sum_{w\in\mathcal{W}} M(w)$. Thus, if $M(W) = \max\{0, \sum_{i=1}^{n} p(t_i) - n + 1\}$, then we are done, since the right-hand side of this formula is the expression for $p^{min}(t_1, \ldots, t_n)$ given in the statement, and it is in every case a lower bound for $p^{min}(t_1, \ldots, t_n)$ (in fact, $p^{min}(t_1, \ldots, t_n)$ can not be less than the case that the tuples are pairwise disconnected in H). Otherwise, it must be the case that $M(W) > \max\{0, \sum_{i=1}^{n} p(t_i) - n + 1\}$. Assume that $\sum_{i=1}^{n} p(t_i) - n + 1 > 0$ (the case that $\max\{0, \sum_{i=1}^{n} p(t_i) - n + 1\} = 0$ can be proved similarly). Hence, we are in the case that $M(W) = \sum_{i=1}^{n} p(t_i) - n + 1 + \epsilon > 0$, with $\epsilon > 0$. Since $M(W') = p(t_n) - M(W)$, this means that $M(W') = p(t_n) - \left(\sum_{i=1}^{n} p(t_i) - n + 1 + \epsilon\right) = -\sum_{i=1}^{n-1} p(t_i) + (n-1) - \epsilon$. From the induction hypothesis, the term $-\sum_{i=1}^{n-1} p(t_i) + (n-1)$ is equal to $1 - p^{min}(t_1, \ldots, t_{n-1})$, thus we have: $M(W') = 1 - p^{min}(t_1, \ldots, t_{n-1}) - \epsilon$. Since $p^{min}(t_1, \ldots, t_{n-1})$ is exactly the overall probability, according to M, of the possible worlds containing all the tuples $t_1, \ldots, t_{n-1}$, we have that $1 - p^{min}(t_1, \ldots, t_{n-1}) = M(W') + M(W''')$, thus we obtain: $M(W') = M(W') + M(W''') - \epsilon$. This means that $M(W''') = \epsilon$, where $\epsilon > 0$. That is, the overall probability of the possible worlds in W''' is equal to the difference ϵ between $M(W)$ and the value $\sum_{i=1}^{n} p(t_i) - n + 1$ that we want to obtain for the cumulative probability of the worlds in W. We now show how M can be modified in order to obtain a model M' such that $M'(W)$ is exactly this value. We construct M' as follows. Let $w'''_1, \ldots, w'''_k$ be the possible worlds in W''' such that $M(w'''_i) > 0$, for each $i \in [1..k]$. Take k values $\epsilon_1, \ldots, \epsilon_k$, where each ϵ_i is equal to $M(w'''_i)$. Hence $\sum_{i=1}^{k} \epsilon_i = \epsilon$. Then, for each $i \in [1..k]$, let $M'(w'''_i) = M(w'''_i) - \epsilon_i = 0$, and, for each $w''' \in W''' \setminus \{w'''_1, \ldots, w'''_k\}$, $M'(w''') = 0$. This way, $M'(W''') = \sum_{w'''\in W'''} M'(w''') = M(W''') - \epsilon = 0$. For each w'''_i (with $i \in [1..k]$), let w'_i be the possible world in W'

"corresponding" to w_i''': that is, w_i' is the possible world $w_i''' \cup \{t_n\}$. The, for each $i \in [1..k]$, let $M'(w_i') = M(w_i') + \epsilon_i$, and, for each $w' \in W' \setminus \{w_1', \ldots, w_k'\}$, $M'(w') = M(w')$. This way, $M'(W') = \sum_{w' \in W'} M'(w') = M(W') + \epsilon$. Basically, we are constructing the model M' by "moving" ϵ from the overall probability assigned by M to the worlds of W''' towards the worlds of W'. Observe that every world $w_i' \in W'$ such that $M'(w_i') > 0$ is consistent w.r.t. $\mathcal{IC}$, for the following reason. If $M'(w_i') = M(w_i')$ the property derives from the fact that M is a model. Otherwise, we are in the case that $w_i' = w_i''' \cup \{t_n\}$, where $M(w_i''') > 0$. Since M is a model, $M(w_i''') > 0$ implies that w_i''' is consistent w.r.t. $\mathcal{IC}$. Then, adding t_n to w_i''' to obtain w_i' has no impact on the consistency: w_i' does not contain at least one tuple among $t_1, \ldots, t_{n-1}$, and from the fact that any hyperedge of $HG(D^p, \mathcal{IC})$ containing t_n contains all the tuples $t_1, \ldots, t_{n-1}$ no constraint encoded by the hyperedges containing t_n is fired in w_i'.

It is easy to see that the strategy that we used to move ϵ from the overall probability of W''' to W' does not change the overall probabilities assigned to the tuples different from t_n in the worlds in $W' \cup W''$, but it changes the overall probability assigned to tuple t_n in the same worlds, as it is increased by ϵ. Hence, to adjust this, we perform an analogous reasoning to "move" ϵ from the overall probability $M(W)$ (which is at least ϵ and whose worlds contain t_n) to the overall probability assigned to W'' (which contains the same worlds of W deprived of t_n). Thus, we define M' by "moving" portions of ϵ from the worlds of W to the corresponding worlds of W'' (where the corresponding worlds are those having the same tuples except from t_n), analogously to what done before from the worlds of W''' to those of W'. This way, we obtain that $M'(W) = M(W) - \epsilon$ and $M'(W'') = M(W) + \epsilon$. Also in this case, M' does not assign a non-zero probability to inconsistent worlds of W'': for any w_i'' such that $M'(w_i'') > M(w_i'')$, it is the case that $M(w_i) > 0$ (where $w_i = w_i'' \cup \{t_n\}$, which means that w_i is consistent, and thus w_i'' (which results from removing a tuple from w_i) must be consistent as well (removing a tuple cannot fire any denial constraint). Finally, observe that this strategy for moving ϵ from the cumulative probability of W to W'' does not alter the marginal probabilities of the tuples different from t_n in these worlds.

Therefore, M' is a model for D^p w.r.t. $\mathcal{IC}$ which assigns to W a cumulative probability equal to $M'(W) = M(W) - \epsilon = \sum_{i=1}^{n} p(t_i) - n + 1$, which ends the proof. □

Appendix A.2. Proof of Theorem 3

In order to prove Theorem 3, we exploit a property that holds for particular conflict hypergraphs, called *chains*. Basically, a chain is the hypergraph resulting from removing a hyperedge from a ring. Thus, a chain consists of a sequence of hyperedges $e_1, \ldots, e_n$ where all and only the pairs of consecutive hyperedges have non-empty intersection (differently from the ring, $e_1 \cap e_n = \emptyset$).

Given a chain $C = e_1, \ldots, e_n$, we say that n is its length, and denote it with $\mathit{length}(C)$. Moreover, for each $i \in [1..n-1]$, we will use the symbol α_i to denote the intersection $e_i \cap e_{i+1}$ of consecutive hyperedges, and, for each $i \in [1..n]$, we will use the symbol β_i to denote $\mathit{ears}(e_i)$, and $\tilde{\beta}_i$ to denote a subset of $\mathit{ears}(e_i)$. Finally, $\mathit{sub}(C)$ will denote the subsequence $e_2, \ldots, e_{n-1}$ of the hyperedges in C.

In the following, given a set of tuples X, we will use the term "*event X*" to denote the event that all the tuples in the set X co-exist. Furthermore, $p_H^{min}(E)$ will denote the minimum probability of the event E involving the tuples of the database D^p when the conflict hypergraph contains only the hyperedges in H.

Lemma 3. *Let D^p be a PDB instance of $\mathcal{D}^p$ such that $D^p \models \mathcal{IC}$. Assume that $HG(D^p, \mathcal{IC})$ is the chain $C = e_1, \ldots, e_n$ (with $n > 1$). Moreover, let $\tilde{\beta}_1, \tilde{\beta}_n$ be subsets of the ears β_1, β_n of e_1 and e_n, respectively. Then:*

$$p_C^{min}(\tilde{\beta}_1 \cup \tilde{\beta}_n) = \max\left\{0,\ p_\emptyset^{min}(\tilde{\beta}_1) + p_\emptyset^{min}(\tilde{\beta}_n) - \left[1 - p_{sub(C)}^{min}(\alpha_1 \cup (\beta_1 \setminus \tilde{\beta}_1) \cup \alpha_{n-1} \cup (\beta_n \setminus \tilde{\beta}_n))\right]\right\}$$

where: $p_{sub(C)}^{min}(\alpha_1 \cup (\beta_1 \setminus \tilde{\beta}_1) \cup \alpha_{n-1} \cup (\beta_n \setminus \tilde{\beta}_n)) = \max\left\{0,\ p_{sub(C)}^{min}(\alpha_1 \cup \alpha_{n-1}) + p_\emptyset^{min}\left((\beta_1 \setminus \tilde{\beta}_1) \cup (\beta_n \setminus \tilde{\beta}_n)\right) - 1\right\}$ *and, for any set of tuples γ,* $p_\emptyset^{min}(\gamma) = \max\left\{0, \sum_{t \in \gamma} p(t) - |\gamma| + 1\right\}$.

Proof. $p(\tilde{\beta}_1 \cup \tilde{\beta}_n)$ can be minimized as follows.
1) We start from any model M of D^p minimizing the portion of the probability space where neither the event $\tilde{\beta}_1$ nor the event $\tilde{\beta}_n$ can occur. That is, M is any model minimizing the probability of the event $E = \alpha_1 \cup (\beta_1 \setminus \tilde{\beta}_1) \cup \alpha_{n-1} \cup (\beta_n \setminus \tilde{\beta}_n)$ (this event is mutually exclusive with both $\tilde{\beta}_1$ and $\tilde{\beta}_n$ due to hyperedges e_1 and e_n). It is easy to see that M is also a model for D^p w.r.t. the conflict hypergraph $\mathit{sub}(C)$, and that the minimum probability $p_{sub(C)}^{min}(E)$ of E w.r.t. $\mathit{sub}(C)$ is equal to the minimum probability $p_C^{min}(E)$ of E w.r.t. C. We denote this probability as Y.

2) We re-distribute the tuples in $\tilde{\beta}_1 \cup \tilde{\beta}_n$ over the portion of size $1-Y$ of the probability space not assigned to E, so that $p(\tilde{\beta}_1) = p^{min}_{\emptyset}(\tilde{\beta}_1)$ and $p(\tilde{\beta}_2) = p^{min}_{\emptyset}(\tilde{\beta}_n)$, and with the aim of minimizing the intersection of the events $\tilde{\beta}_1$ and $\tilde{\beta}_n$. The fact that the events $\tilde{\beta}_1$ and $\tilde{\beta}_n$ can be simultaneously assigned their minimum probabilities $p^{min}_{\emptyset}(\tilde{\beta}_1)$ and $p^{min}_{\emptyset}(\tilde{\beta}_n)$, respectively, derives from Lemma 1 and from the consistency of D^p w.r.t. C. This yields a (possibly) new model M' for D^p w.r.t. the "original" chain C where $p(\tilde{\beta}_1 \cup \tilde{\beta}_n) = \max\left\{0,\ p^{min}(\tilde{\beta}_1) + p^{min}(\tilde{\beta}_n) - [1-Y]\right\}$. In fact, viewing the available probability space as a segment of length $1-Y$, this corresponds to assigning the left-most part of the segment of length $p^{min}(\tilde{\beta}_1)$ to event $\tilde{\beta}_1$, and the right-most part of length $p^{min}(\tilde{\beta}_n)$ to event $\tilde{\beta}_n$. This way, the probability of the intersection is the length of the segment portion (if any) assigned to both $\tilde{\beta}_1$ and $\tilde{\beta}_n$. In brief, we obtain the formula reported in the statement for $p^{min}(\tilde{\beta}_1 \cup \tilde{\beta}_n)$.

The formula for $p^{min}(\alpha_1 \cup (\beta_1 \setminus \tilde{\beta}_1) \cup \alpha_{n-1} \cup (\beta_n \setminus \tilde{\beta}_n))$ can be proved with an analogous reasoning, while the formula for $p^{min}_{\emptyset}(\gamma)$ follows from Lemma 1. □

Theorem 3. *Given an instance D^p of $\mathcal{D}^p$, if $HG(D^p, \mathcal{IC}) = \langle N, E\rangle$ is a ring, then $D^p \models \mathcal{IC}$ iff both the following hold: 1) $\forall e \in E,\ \sum_{t\in e} p(t) \leq |e| - 1$; 2) $\sum_{t\in N} p(t) - |N| + \lceil \frac{|E|}{2} \rceil \leq 0$.*

Proof. In the following, we will denote the ring $HG(D^p, \mathcal{IC})$ as $\mathcal{R} = e_1, \ldots, e_n, e_{n+1}$, and, for each $i \in [1..n+1]$, the ears of e_i as ε_i, and, for each $i \in [1..n]$, the intersection $e_i \cap e_{i+1}$ as γ_i, and $e_1 \cap e_{n+1}$ as γ_0. Moreover, we will denote as $C = e_1, \ldots, e_n$ the chain obtained from ring $\mathcal{R}$ by removing the edge e_{n+1}. We now prove the left-to-right and right-to-left implications separately.

($\Rightarrow$): We first show that, if $D^p \models \mathcal{IC}$ and $HG(D^p, \mathcal{IC})$ is a ring, then both Condition 1. and 2. hold. Condition 1. trivially follows from the fact that the proof of the left-to-right implication of Theorem 2 holds for general conflict hypergraphs.

We now focus on Condition 2. As D^p is consistent w.r.t. $\mathcal{R}$, the presence of hyperedge e_{n+1} in $HG(D^p, \mathcal{IC})$ implies that the minimum probability that the tuples in e_{n+1} co-exist is equal to 0. That is, $p^{min}_{\mathcal{R}}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) = 0$. On the other hand, $p^{min}_{C}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) \leq p^{min}_{\mathcal{R}}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1})$, thus it must hold that $p^{min}_{C}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) = 0$. Since, according to the conflict hypergraph C, no correlation is imposed between the events $(\gamma_0 \cup \gamma_n)$ and ε_{n+1}, we also have that $p^{min}_{C}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) = \max\{0, p^{min}_{C}(\gamma_0 \cup \gamma_n) + p^{min}_{\emptyset}(\varepsilon_{n+1}) - 1\}$ (see Fact 2). Hence, the following inequality must hold:

$$p^{min}_{C}(\gamma_0 \cup \gamma_n) + p^{min}_{\emptyset}(\varepsilon_{n+1}) - 1 \leq 0. \tag{A.1}$$

We now show that inequality (A.1) entails that Condition 2. holds. First, observe that γ_0 and γ_n are subsets of the ears of e_1 and e_n, respectively, w.r.t. the hypergraph C. Hence, since C is a chain, we can apply Lemma 3 to obtain $p^{min}_{C}(\gamma_0 \cup \gamma_n)$ in function of $p^{min}_{sub(C)}(\gamma_1 \cup \gamma_{n-1})$. Thus, by recursively applying ($\lfloor \frac{n}{2} \rfloor$ times) Lemma 3, we obtain the following expression for $p^{min}_{C}(\gamma_0 \cup \gamma_n)$ (where $x = \lfloor \frac{n}{2} \rfloor - 1$ and $y = \lceil \frac{n}{2} \rceil + 1$):

$$\begin{aligned}
&\max\left\{0, \max\left\{0, \textstyle\sum_{t\in\gamma_0} p(t) - |\gamma_0| + 1\right\} + \max\left\{0, \textstyle\sum_{t\in\gamma_n} p(t) - |\gamma_n| + 1\right\} - 1 + \right.\\
&\qquad \max\left\{0, \max\left\{0, \max\left\{0, \textstyle\sum_{t\in\gamma_1} p(t) - |\gamma_1| + 1\right\} + \max\left\{0, \textstyle\sum_{t\in\gamma_{n-1}} p(t) - |\gamma_{n-1}| + 1\right\} - 1 + \right.\right.\\
&\qquad\qquad \cdots\\
&\qquad\qquad\quad \max\left\{0, \max\left\{0, \textstyle\sum_{t\in\gamma_x} p(t) - |\gamma_x| + 1\right\} + \max\left\{0, \textstyle\sum_{t\in\gamma_y} p(t) - |\gamma_y| + 1\right\} - 1 + P\right\} + \\
&\qquad\qquad \cdots\\
&\qquad\quad \left.\left. \max\left\{0, \textstyle\sum_{t\in(\varepsilon_2\cup\varepsilon_{n-1})} p(t) - |\varepsilon_2 \cup \varepsilon_{n-1}| + 1\right\} - 1\right.\\
&\qquad \}+\\
&\quad \max\left\{0, \textstyle\sum_{t\in(\varepsilon_1\cup\varepsilon_n)} p(t) - |\varepsilon_1 \cup \varepsilon_n| + 1\right\} - 1\\
&\}
\end{aligned}$$

where:

$$P = \begin{cases} p^{min}_{\emptyset}(\gamma_{x+1}) & \text{if } n \text{ is even;} \\ p^{min}_{e_{y-1}}(\gamma_{x+1} \cup \gamma_{y-1}) & \text{otherwise.} \end{cases}$$

In this formula, $p^{min}_{\emptyset}(\gamma_{x+1}) = \max\{0, \sum_{t\in\gamma_{x+1}} p(t) - |\gamma_{x+1}| + 1\}$, and $p^{min}_{e_{y-1}}(\gamma_{x+1} \cup \gamma_{y-1}) = \max\{0, \sum_{t\in(\gamma_{x+1}\cup\gamma_{y-1})} p(t) - |(\gamma_{x+1} \cup \gamma_{y-1})| + 1\}$ (the latter follows from applying Lemma 1).

The value of $p_C^{min}(\gamma_0 \cup \gamma_n)$ is greater than or equal to the sum S of the non-zero terms that occur in the expression obtained so far, that is:

$$S = \begin{cases} \sum_{t\in(N\setminus\varepsilon_{n+1})} p(t) - (|N| - |\varepsilon_{n+1}|) + \frac{n}{2} + 1, \\ \text{if the length } n \text{ of the chain } C \text{ is even;} \\ \\ \sum_{t\in(N\setminus\varepsilon_{n+1})} p(t) - (|N| - |\varepsilon_{n+1}| - |\varepsilon_{x+2}|) + \lfloor \frac{n}{2} \rfloor + 1, \\ \text{if the length } n \text{ of } C \text{ is odd.} \end{cases}$$

The fact that $p_C^{min}(\gamma_0 \cup \gamma_n) \geq S$ straightforwardly follows from that S is obtained by summing also possibly negative contributions of terms of the form $p_\emptyset^{min}(Z) = \sum_{t\in Z} p(t) - |Z| + 1$, which are not considered when evaluating $p_C^{min}(\gamma_0 \cup \gamma_n)$, since invocations of the max function return non-negative values only.

As the number of edges in the ring $\mathcal{R}$ is $|E| = n + 1$, the value of S is in every case greater than or equal to

$$S' = \sum_{t\in(N\setminus\varepsilon_{n+1})} p(t) - (|N| - |\varepsilon_{n+1}|) + \left\lceil \frac{|E|}{2} \right\rceil$$

In brief, we have obtained $S' \leq S \leq p_C^{min}(\gamma_0 \cup \gamma_n)$.

Since $D^p \models \mathcal{IC}$ implies that $p_C^{min}(\gamma_0 \cup \gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 1 \leq 0$ (equation (A.1)), we obtain $S' + p_\emptyset^{min}(\varepsilon_{n+1}) - 1 \leq 0$. By replacing S' and $p_\emptyset^{min}(\varepsilon_{n+1})$ with the corresponding formulas, we obtain

$$\sum_{t\in(N\setminus\varepsilon_{n+1})} p(t) - (|N| - |\varepsilon_{n+1}|) + \lceil \frac{|E|}{2} \rceil + \sum_{t\in\varepsilon_{n+1}} p(t) - |\varepsilon_{n+1}| \leq 0$$

that is, $\sum_{t\in N} p(t) - |N| + \left\lceil \frac{|E|}{2} \right\rceil \leq 0$.

($\Leftarrow$): We now prove the right-to-left implication, reasoning by contradiction. Assume that both Condition 1. and 2. hold, but D^p is not consistent w.r.t. the conflict hypergraph $\mathcal{R}$. However, since C is a hypertree and Condition 1. holds, from Theorem 2 we have that D^p is consistent w.r.t. the conflict hypergraph C. In particular, it must be the case that $p_C^{min}(e_{n+1}) = p_C^{min}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) > 0$: otherwise, any model of D^p w.r.t. C assigning probability 0 to the event $(\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}$ would be also a model for D^p w.r.t. $\mathcal{R}$, which is in contrast with the contradiction hypothesis.

Since, according to the conflict hypergraph C, no correlation is imposed between the events $(\gamma_0 \cup \gamma_n)$ and ε_{n+1}, we also have that $p_C^{min}((\gamma_0 \cup \gamma_n) \cup \varepsilon_{n+1}) = \max\{0, p_C^{min}(\gamma_0 \cup \gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 1\}$ (see Fact 2). Hence, the following inequality must hold:

$$p_C^{min}(\gamma_0 \cup \gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 1 > 0 \tag{A.2}$$

which also implies both $p_C^{min}(\gamma_0 \cup \gamma_n) > 0$ and $p_\emptyset^{min}(\varepsilon_{n+1}) > 0$ (as probabilities values are bounded by 1).

By applying Lemma 4, we obtain that $p_C^{min}(\gamma_0 \cup \gamma_n)$ is equal to

$$\max\{0, p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) - 1 + \max\{0, p_{sub(C)}^{min}(\gamma_1 \cup \gamma_{n-1}) + p_\emptyset^{min}(\varepsilon_1 \cup \varepsilon_n) - 1\}\}$$

As shown above, $p_C^{min}(\gamma_0 \cup \gamma_n) > 0$, thus the expression for $p_C^{min}(\gamma_0 \cup \gamma_n)$ can be simplified into:

$$p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) - 1 + \max\{0, p_{sub(C)}^{min}(\gamma_1 \cup \gamma_{n-1}) + p_\emptyset^{min}(\varepsilon_1 \cup \varepsilon_n) - 1\}$$

By replacing $p_C^{min}(\gamma_0 \cup \gamma_n)$ with this formula in equation (A.2), we obtain

$$p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 2 + \max\{0, p_{sub(C)}^{min}(\gamma_1 \cup \gamma_{n-1}) + p_\emptyset^{min}(\varepsilon_1 \cup \varepsilon_n) - 1\} > 0 \tag{A.3}$$

Since $p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 2 \leq p_\emptyset^{min}(\gamma_0 \cup \gamma_n \cup \varepsilon_{n+1})$ (which follows from applying twice Fact 2), and $p_\emptyset^{min}(\gamma_0 \cup \gamma_n \cup \varepsilon_{n+1}) = \max\{0, \sum_{t\in(\gamma_0\cup\gamma_n\cup\varepsilon_{n+1})} p(t) - |(\gamma_0 \cup \gamma_n \cup \varepsilon_{n+1})| + 1\}$, and $\sum_{t\in(\gamma_0\cup\gamma_n\cup\varepsilon_{n+1})} p(t) - |(\gamma_0 \cup \gamma_n \cup \varepsilon_{n+1})| + 1 \leq 0$ (Condition 1. over hyperedge e_{n+1}), we obtain that $p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 2 \leq 0$. Hence, the second argument of max in equation (A.3) must be strictly positive, thus equation (A.3) can be rewritten as:

$$p_\emptyset^{min}(\gamma_0) + p_\emptyset^{min}(\gamma_n) + p_\emptyset^{min}(\varepsilon_{n+1}) - 2 + p_{sub(C)}^{min}(\gamma_1 \cup \gamma_{n-1}) + p_\emptyset^{min}(\varepsilon_1 \cup \varepsilon_n) - 1 > 0 \tag{A.4}$$

where $p^{min}_{sub(C)}(\gamma_1 \cup \gamma_{n-1}) > 0$ and $p^{min}_{\emptyset}(\varepsilon_1 \cup \varepsilon_n) > 0$ (otherwise, the second argument of max in equation (A.3) could not be strictly positive, being probability values bounded by 1).

Observe that all the terms of the form p^{min} occurring in (A.4) are strictly positive. In fact, we have already shown that this holds for $p^{min}_{\emptyset}(\varepsilon_{n+1})$, $p^{min}_{sub(C)}(\gamma_1 \cup \gamma_{n-1})$, and $p^{min}_{\emptyset}(\varepsilon_1 \cup \varepsilon_n)$. As regards $p^{min}_{\emptyset}(\gamma_0)$, the fact that it is strictly greater than 0 derives from the $p^{min}_{\emptyset}(\gamma_0) = p^{min}_{C}(\gamma_0)$ (which is due to Lemma 1, as γ_0 is a matryoshka w.r.t. C), and $p^{min}_{C}(\gamma_0) \geq p^{min}_{C}(\gamma_0 \cup \gamma_n)$, where $p^{min}_{C}(\gamma_0 \cup \gamma_n) > 0$, as shown before. The same reasoning suffices to prove that $p^{min}_{\emptyset}(\gamma_n) > 0$.

The fact that all the terms of the form $p^{min}_{\emptyset}$ in (A.4) are strictly positive implies that we can replace them with the corresponding formulas given in Lemma 1, simplified by eliminating the max operator. Therefore, we obtain:

$$\begin{aligned}&\left(\textstyle\sum_{t\in\gamma_0} p(t) - |\gamma_0| + 1\right) + \left(\textstyle\sum_{t\in\gamma_n} p(t) - |\gamma_n| + 1\right) + \left(\textstyle\sum_{t\in\varepsilon_{n+1}} p(t) - |\varepsilon_{n+1}| + 1\right) + \ p^{min}_{sub(C)}(\gamma_1 \cup \gamma_{n-1}) + \\ &+ \left(\textstyle\sum_{t\in\varepsilon_1} p(t) - |\varepsilon_1| + 1\right) + \left(\textstyle\sum_{t\in\varepsilon_n} p(t) - |\varepsilon_n| + 1\right) - 1 - 3 > 0\end{aligned} \tag{A.5}$$

By recursively applying the same reasoning on $p^{min}_{sub(C)}(\gamma_1 \cup \gamma_{n-1})$ a number of times equal to $\lfloor \frac{n}{2} \rfloor$, the term on the left-hand side of equation (A.5) can be shown to be less than or equal to $\sum_{t\in N} p(t) - |N| + \left\lceil \frac{|E|}{2} \right\rceil$ (depending on whether n is even or not, analogously to the proof of the inverse implication). Thus, we obtain $\sum_{t\in N} p(t) - |N| + \left\lceil \frac{|E|}{2} \right\rceil > 0$, which contradicts Condition 2. □

Appendix A.3. Proofs of theorems 4, 5, 6, 7, and 8

Theorem 4. *If $\mathcal{IC}$ consists of a join-free denial constraint, then* cc *is in PTIME. In particular, $D^p \models \mathcal{IC}$ iff, for each hyperedge e of $HG(D^p, \mathcal{IC})$, it holds that $\sum_{t\in e} p(t) \leq |e| - 1$.*

Proof. Let $\mathcal{IC}$ consist of the denial constraint *ic* having the form: $\neg[R_1(\vec{x}_1) \wedge \cdots \wedge R_m(\vec{x}_m) \wedge \phi_1(\vec{x}_1) \wedge \cdots \wedge \phi_m(\vec{x}_m)]$, where no variable occurs in two distinct relation atoms of *ic*, and, for each built-in predicate occurring in $\phi_1(\vec{x}_1) \wedge \cdots \wedge \phi_m(\vec{x}_m)$ at least one term is a constant. Given an instance D^p of $\mathcal{D}^p$, we show that $D^p \models \mathcal{IC}$ iff for each hyperedge e of $HG(D^p, \mathcal{IC})$, it holds that $\sum_{t\in e} p(t) \leq |e| - 1$.

($\Rightarrow$): It straightforwardly follows for the fact that, as pointed out in the core of the paper after Theorem 2, the condition that, for each hyperedge e of $HG(D^p, \mathcal{IC})$, $\sum_{t\in e} p(t) \leq |e| - 1$ is a necessary condition for the consistency in the presence of any conflict hypergraph.

($\Leftarrow$): For each $i \in [1..m]$, let R_{ϕ_i} be the maximal set of tuples in the instance of R_i such that every tuple $t_i \in R_{\phi_i}$ satisfies $R_i(\vec{x}_i) \wedge \phi_i(\vec{x}_i)$.

It is easy to see that $HG(D^p, \mathcal{IC})$ consists of the set of hyperedges $\left\{\{t_1, \ldots, t_m\} \mid \forall i \in [1..m]\ t_i \in R_{\phi_i}\right\}$. Observe that not all the hyperdeges in $HG(D^p, \mathcal{IC})$ have size m, as the same relation scheme may appear several times in *ic*. That is, in the case that there are $i, j \in [1..m]$ with $i < j$ such that $R_{\phi_i} \cap R_{\phi_j} \neq \emptyset$, the tuples t_i and t_j occurring in the same hyperedge $\{t_1, \ldots, t_i, \ldots, t_j, \ldots, t_m\}$ may coincide, thus this hyperedge has size less than m.

From the hypothesis, it holds that, for every hyperedge e of $HG(D^p, \mathcal{IC})$, it must be the case that $\sum_{t\in e} p(t) \leq |e| - 1$. Let e^* be the hyperedge in $HG(D^p, \mathcal{IC})$ such that $|e| - 1 - \sum_{t\in e} p(t)$ is the minimum, that is,

$$e^* = argmin_{e\in HG(D^p,\mathcal{IC})}\left(|e| - 1 - \sum_{t\in e} p(t)\right).$$

For the sake of simplicity of presentation we consider the case that e^* has size m, and denote its tuples as $t_1, \ldots, t_m$. The generalization to the case that the size of e^* is less than m is straightforward.

Let S be a subset of D^p. We denote with D^p_S the subset of D^p containing only the tuples in S. Let Pr_{e^*} be a model in $\mathcal{M}(D^p_{e^*}, \mathcal{IC})$. Moreover, let $t'_1, \ldots, t'_n$ be the tuples in D^p / e^*.

In the following, we will define a sequence of interpretations $Pr_0, Pr_1, \ldots, Pr_n$ such that, for each $i \in [0..n]$, Pr_i is a model in $\mathcal{M}(D^p_{e^*\cup\{t'_j | j\leq i\}}, \mathcal{IC})$.

We start by taking Pr_0 equal to Pr_{e^*}. At the i^{th} step we consider tuple t'_i and define Pr_i as follows:

1. In the case that, for each $j \in [1..m]$, it holds that $t'_i \notin R_{\phi_j}$, we define, for each possible world w in $pwd(e^* \cup \{t'_j | j \leq i\})$, $Pr_i(w) = Pr_{i-1}(w \setminus \{t'_i\}) \cdot p(t'_i)$, if $t'_i \in w$, and $Pr_i(w) = Pr_{i-1}(w \setminus \{t'_i\}) \cdot (1 - p(t'_i))$, otherwise.

2. Otherwise, if there is $j \in [1..m]$ such that $t'_i \in R_{\phi_j}$, we consider the set J of all the indexes $j \in [1..m]$ such that $t'_i \in R_{\phi_j}$. Moreover, we denote with p_J the sum of the probabilities (computed according to Pr_{i-1}) of all the possible worlds $w \in pwd(e^* \cup \{t'_j | j \leq i-1\})$ such that, for each $j \in J$, the corresponding tuple t_j appearing in e^* belongs also to w, i.e., $p_J = \sum_{w \in pwd(e^* \cup \{t'_j | j \leq i-1\}), s.t. \forall j \in J\, t_j \in w} Pr_{i-1}(w)$.
Then, for each possible world w in $pwd(e^* \cup \{t'_j | j \leq i\})$, we define Pr_i as follows:

 - $Pr_i(w) = Pr_{i-1}(w - \{t'_i\}) \cdot \frac{p(t'_i)}{p_J}$, if $t'_i \in w$ and for each $j \in J$ it holds that $t_j \in w$,
 - $Pr_i(w) = Pr_{i-1}(w - \{t'_i\}) \cdot \frac{max(0, p_J - p(t'_i))}{p_J}$, if $t'_i \notin w$ and for each $j \in J$ it holds that $t_j \in w$,
 - $Pr_i(w) = Pr_{i-1}(w)$, if $t'_i \notin w$ and there is a $j \in J$ such that $t_j \notin w$,
 - $Pr_i(w) = 0$, otherwise.

We prove that for each $i \in [0..n]$ it holds that Pr_i is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i\}}, \mathcal{IC})$ reasoning by induction on i. The proof is straightforward for $i = 0$. We now prove the induction step, that is, we assume that Pr_{i-1} is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i-1\}}, \mathcal{IC})$ and prove that Pr_i is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i\}}, \mathcal{IC})$.

As regards the first case of the definition of Pr_i from Pr_{i-1}, it is easy to see that Pr_i is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i\}}, \mathcal{IC})$ since Pr_i consists in a trivial extension of Pr_{i-1} which takes into account a tuple not correlated with the other tuples in the database.

As regards the second case of the definition of Pr_i from Pr_{i-1}, it is easy to see that, if $p_J \geq p(t'_i)$ than Pr_i guarantees that the condition about the marginal probabilities of all the tuples in $e^* \cup \{t'_j | j \leq i\}$ holds. Moreover, Pr_j assigns zero probability to each possible world w such that $w \not\models \mathcal{IC}$, since, for each possible world w in $pwd(e^* \cup \{t'_j | j \leq i\})$, there is no subset S of w such that for each $i \in [1..m]$ there is a tuple $t \in S$ such that $t \in R_{\phi_i}$. The latter follows from the induction hypothesis, which ensures that Pr_{i-1} is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i-1\}}, \mathcal{IC})$, and from the fact that Pr_i assigns non-zero probability to a possible world w in $pwd(e^* \cup \{t'_j | j \leq i\})$ containing t'_i iff for each $j \in J$ it holds that $t_j \in w$. Specifically, it can not be the case that w contains, for each $x \in [1..m]$ such that $x \notin J$ a tuple $t_x \in R_{\phi_i}$, as otherwise $w - \{t'_i\}$ would satisfy all the conditions expressed in ic, and $w - \{t'_i\}$ would be assigned a non-zero probability by Pr_{i-1}, thus contradicting the induction hypothesis that Pr_{i-1} is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i-1\}}, \mathcal{IC})$.

We now prove that $p_J \geq p(t'_i)$. Reasoning by contradiction, assume that $p_J < p(t'_i)$. From the definition of p_J it follows that $p_J \geq p^{min}(\wedge_{j \in J} t_j)$. Therefore, since $p^{min}(\wedge_{j \in J} t_j)$ is equal to $\max\left\{0, \sum_{j \in J} p(t_j) - |J| + 1\right\}$ it follows that $p(t'_i) > \sum_{j \in J} p(t_j) - |J| + 1$. Consider the hyperdege $e = \{t_x | t_x \in e^* \wedge x \notin J\} \cup \{t'_i\}$. From the definition of e^* it follows that $|e| - 1 - \sum_{t \in e} p(t) \geq |e^*| - 1 - \sum_{t \in e^*} p(t)$. The latter implies that $1 - p(t'_i) \geq |J| - \sum_{j \in J} p(t_j)$, from which it follows that $\sum_{j \in J} p(t_j) - |J| + 1 \geq p(t'_i)$ which is a contradiction. Hence, we can conclude that, in this case Pr_i is a model in $\mathcal{M}(D^p_{e^* \cup \{t'_j | j \leq i\}}, \mathcal{IC})$.

This conclude the proof, as Pr_n is a model in $\mathcal{M}(D^p, \mathcal{IC})$ and then $D^p \models \mathcal{IC}$. □

Theorem 5. *There is an $\mathcal{IC}$ consisting of a non-join-free denial constraint of arity 3 such that* CC *is NP-hard.*

Proof. The reader is kindly requested to read this proof after that of Theorem 7, as the construction used there will be exploited in the reasoning used below.

We show that the reduction from 3-COLORING to CC presented in the hardness proof of Theorem 7 can be rewritten to obtain an instance of CC where $\mathcal{IC}$ contains only a denial constraints having arity equal to 3.

Let $G = \langle N, E \rangle$ be a 3-COLORING instance. We construct an equivalent instance $\langle \mathcal{D}^p, \mathcal{IC}, D^p \rangle$ of CC as follows:

– $\mathcal{D}^p$ consists of the probabilistic relation schemas R^p_1(*Node, Color, P*) and R^p_2(*Node1, Node2, Color1, Color2, P*);
– D^p is the instance of $\mathcal{D}^p$ consisting of the instances r^p_1 of R^p_1, and r^p_2 of R^p_2, defined as follows:

 - for each node $n \in N$, and for each color $c \in$ {*Red, Green, Blue*}, r^p_1 contains the tuple $(n, c, \frac{1}{3})$;
 - for each edge $\{n_1, n_2\} \in E$, and for each color $c \in$ {*Red, Green, Blue*}, r^p_2 contains the tuple $(n_1, n_2, c, c, 1)$; moreover, for each node $n \in N$, and for each pair of distinct colors $c_1, c_2 \in$ {*Red, Green, Blue*}, r^p_2 contains the tuple $(n, n, c_1, c_2, 1)$;

– $\mathcal{IC}$ is the set of denial constraints over $\mathcal{D}^p$ consisting of the constraint: $\neg[R_1(x_1, x_2) \wedge R_1(x_3, x_4) \wedge R_2(x_1, x_3, x_2, x_4)]$.

Basically, the constraint in $\mathcal{IC}$ imposes that adjacent nodes can not be assigned the same color, and the same node can not be assigned more than one color.

Let $\langle \overline{\mathcal{D}}^p, \overline{\mathcal{IC}}, \overline{D}^p \rangle$ be the instance of cc defined in the hardness proof of Theorem 1, where it was shown that an instance G of 3-COLORING is 3-colorable iff $\overline{D}^p \models \overline{\mathcal{IC}}$. It is easy to see that $D^p \models \mathcal{IC}$ iff $\overline{D}^p \models \overline{\mathcal{IC}}$, which completes the proof. □

Theorem 6. *If $\mathcal{IC}$ consists of a BEGD, then* cc *is in PTIME.*

Proof. Let the BEGD in $\mathcal{IC}$ be:

$$ic = \neg[R_1(\vec{x}, \vec{y}_1) \wedge R_2(\vec{x}, \vec{y}_2) \wedge z_1 \neq z_2],$$

where each z_i (with $i \in \{1, 2\}$) is a variable in $\vec{y}_1$ or $\vec{y}_2$. That is, for the sake of presentation, we assume that the conjunction of built-in predicates in ic consists of one conjunct only (this yields no loss of generality, as it is easy to see that the reasoning used in the proof is still valid in the presence of more conjuncts). We consider two cases separately.

Case 1: $R_1 = R_2$, that is, only one relation name occurs in ic. Let $\overline{X}$ be the set of attributes in $Attr(R_1)$ corresponding to the variables in $\vec{x}$, and let Z_1 and Z_2 be the attributes in $Attr(R_1)$ corresponding to the variables z_1 and z_2, respectively. Let r be an instance of R_1.

It is easy to see that the conflict hypergraph $HG(r, \mathcal{IC})$ is a graph having the following structure: for any pair of tuples t_1, t_2, there is the edge (t_1, t_2) in $HG(r, ic)$ iff: 1) $\forall X \in \overline{X}$, $t_1[X] = t_2[X]$, and 2) $t_1[Z_1] \neq t_2[Z_2]$.

This structure of the conflict hypergraph implies a partition of the tuples of r, where the tuples in each set of the partition share the same values of the attributes in $\overline{X}$. Obviously, cc can be decided by considering these sets separately.

For each set G of this partition, we reason as follows. Let $\mathcal{P}_G$ be the set of pairs of values $\langle c_1, c_2 \rangle$ occurring as values of attributes Z_1 and Z_2 in at least one tuple of r (that is, $\mathcal{P}_G$ is the projection of r over Z_1 and Z_2). For each pair $\langle c_1, c_2 \rangle \in \mathcal{P}_G$, let $T[c_1, c_2]$ be the set of tuples in G such that, $\forall t \in T[c_1, c_2]$, $t[Z_1] = c_1$ and $t[Z_2] = c_2$. A first necessary condition for consistency is that there is no pair $\langle c_1, c_2 \rangle \in \mathcal{P}_G$ such that $c_1 \neq c_2$: otherwise, any tuple in $T[c_1, c_2]$ would not satisfy the constraint, thus it would not be possible to put it in any possible world with non-zero probability[2]. Straightforwardly, this condition is also sufficient if z_1 and z_2 belong to the same relation atom. Thus, in this case, the proof ends, as checking this condition can be done in polynomial time.

Otherwise, if z_1 and z_2 belong to different relation atoms and if the above-introduced necessary condition holds, we proceed as follows. From what said above, it must be the case that $\mathcal{P}_G$ contains only pairs of the form $\langle c, c \rangle$, and, correspondingly, all the sets $T[c_1, c_2]$ are of the form $T[c, c]$. For each $T[c, c]$, let $\widetilde{p}(T[c, c])$ be the maximum probability of the tuples in $T[c, c]$ (i.e., $\widetilde{p}(T[c, c]) = \max_{t \in T[c,c]}\{p(t)\}$. Moreover, for each $\langle c, c \rangle \in \mathcal{P}_G$, take the tuple t_c in G such that $p(t_c) = \widetilde{p}(T[c, c])$, and let $\mathcal{T}_G$ be the set of these tuples. We show that cc is *true* iff, for each G, the following inequality (which can be checked in polynomial time) holds:

$$\sum_{\langle c,c \rangle \in \mathcal{P}_G} \widetilde{p}(T[c, c]) \leq 1 \tag{A.6}$$

($\Rightarrow$): Reasoning by contradiction, assume that, for a group G, inequality (A.6) does not hold, but there is a model for the PDB w.r.t. $\mathcal{IC}$.

The constraint entails that, for each pair of distinct tuples $t_1, t_2 \in \mathcal{T}_G$, there is the edge (t_1, t_2) in $HG(r, \mathcal{IC})$. Hence, there is a clique in $HG(r, ic)$ consisting of the tuples in $\mathcal{T}_G$. Since the sum of the probabilities of the tuples in $\mathcal{T}_G$ is greater than 1 (by contradiction hypothesis), and since cc is *true* only if, for each clique in the conflict hypergraph, the sum of the probabilities in the clique does not exceed 1, it follows that cc is *false*.

($\Leftarrow$): It is straightforward to see that there is model for $\mathcal{T}_G$ w.r.t. ic, since the sum of the probabilities of the tuples in $\mathcal{T}_G$ is less than or equal to 1, and since the tuples in $\mathcal{T}_G$ describe a clique in $HG(\mathcal{T}_G, ic)$. Since, for each $\langle c, c \rangle \in \mathcal{P}_G$,

[2] Obviously, we assume that there is no tuple with zero probability, as tuples with zero probability can be discarded from the database instance.

the tuple t_c in $\mathcal{T}_G$ is such that its probability is not less than the probability of every other tuple in $T[c,c]$, it is easy to see that a model M for G w.r.t. ic can be obtained by putting the tuples in $T[c,c]$ other than t_c in the portion of the probability space occupied by the worlds containing t_c.

Case 2: $R_1 \neq R_2$. We assume that $z_1 \in \vec{y_1}$ and $z_2 \in \vec{y_2}$, that is, two distinct relation names occur in ic, and the variables of the inequality predicate belongs to different relation atoms. In fact, the case that z_1 and z_2 belong to the same relation atom can be proved by reasoning analogously.

Let $\overline{X}_1$ and $\overline{X}_2$ be, respectively, the set of attributes in $Attr(R_1)$ and $Attr(R_2)$ corresponding to the variables in $\vec{x}$, and let Z_1 and Z_2 be the attributes in $Attr(R_1)$ and $Attr(R_2)$ corresponding to the variables z_1 and z_2, respectively. Let r_1 be the instance of R_1, and r_2 be the instance of R_2.

Observe that ic does not impose any condition between pairs of tuples $t_1 \in r_1$ and $t_2 \in r_2$ such that there are attributes $X_1 \in \overline{X}_1$ and $X_2 \in \overline{X}_2$ such that $t_1[X_1] \neq t_2[X_2]$. This entails that cc can be decided by considering the consistency of the tuples of r_1 and r_2 sharing the same combination of values for the attributes corresponding to the variables in $\vec{x}$ separately from the tuples sharing different combinations of values for the same attributes. For each combination $\vec{v} = v_1, \ldots, v_k$ of values for these attributes (i.e., $\forall \vec{v} \in \Pi_{\overline{X}_1}(r_1) \cap \Pi_{\overline{X}_2}(r_2)$), let $G_1(\vec{v})$ and $G_2(\vec{v})$ be the sets of tuples of r_1 and r_2, respectively, where the attributes corresponding to the variables in $\vec{x}$ have values $v_1, \ldots, v_k$. Let $\mathcal{V}(G_1(\vec{v})) = \{t[Z_1] \mid t \in G_1(\vec{v})\}$ and $\mathcal{V}(G_2(\vec{v})) = \{t[Z_2] \mid t \in G_2(\vec{v})\}$. For each $c_1 \in \mathcal{V}(G_1(\vec{v}))$ (resp., $c_2 \in \mathcal{V}(G_2(\vec{v}))$), let $T_1[c_1]$ (resp., $T_2[c_2]$) be the set of tuples t of $G_1(\vec{v})$ (resp., $G_2(\vec{v})$) such that $t[Z_1] = c_1$ (resp., $t[Z_2] = c_2$). Moreover, for each $c_1 \in \mathcal{V}(G_1(\vec{v}))$ (resp., $c_2 \in \mathcal{V}(G_2(\vec{v}))$), let $\widetilde{p}(T_1[c_1])$ (resp., $\widetilde{p}(T_2[c_2])$) be the maximum probability of the tuples in $T_1[c_1]$ (resp., $T_2[c_2]$).

We show that cc is *true* iff, $\forall \vec{v} \in \Pi_{\overline{X}_1}(r_1) \cap \Pi_{\overline{X}_2}(r_2)$, it is the case that:

$$\forall c_1 \in \mathcal{V}(G_1(\vec{v})) \; \forall c_2 \in \mathcal{V}(G_2(\vec{v})) \text{ s.t. } c_1 \neq c_2, \text{ it holds that } \widetilde{p}(T_1[c_1]) + \widetilde{p}(T_2[c_2]) \leq 1 \tag{A.7}$$

($\Rightarrow$): Reasoning by contradiction, assume that the database is consistent but there are $c_1 \in \mathcal{V}(G_1(\vec{v}))$ and $c_2 \in \mathcal{V}(G_2(\vec{v}))$, with $c_1 \neq c_2$, such that $\widetilde{p}(T_1[c_1]) + \widetilde{p}(T_2[c_2]) > 1$. Hence, there are tuples $t_1 \in T_1[c_1]$ and $t_2 \in T_2[c_2]$ such that $p(t_1) + p(t_2) > 1$. As these tuples form a conflicting set, the conflict hypergraph $HG(D^p, \mathcal{IC})$ contains the edge (t_1, t_2). It follows that the condition of Theorem 2, that is a necessary condition for the consistency in the presence of any hypergraph (as pointed out in the core of the paper after Theorem 2), is not satisfied, thus contradicting the hypothesis.

($\Leftarrow$): It suffices to separately consider each $\vec{v} \in \Pi_{\overline{X}_1}(r_1) \cap \Pi_{\overline{X}_2}(r_2)$, and to show that the fact that (A.7) holds for this $\vec{v}$ implies the consistency of the tuples in $G_1(\vec{v}) \cup G_2(\vec{v})$ (as explained above, the consistency can be checked by separately considering the various combinations in $\Pi_{\overline{X}_1}(r_1) \cap \Pi_{\overline{X}_2}(r_2)$).

Let $\widetilde{t_1} \in G_1(\vec{v})$ and $\widetilde{t_2} \in G_2(\vec{v})$ be such that

(i) $\widetilde{t_1} \in T_1[c_1]$ and $\widetilde{t_2} \in T_2[c_2]$, with $c_1 \neq c_2$; and

(ii) among the pair of tuples satisfying the above conditions, $\widetilde{t_1}$ and $\widetilde{t_2}$ have maximum probability w.r.t. the tuples in $G_1(\vec{v})$ and $G_2(\vec{v})$, respectively.

If these two tuples do not exist, it means that the set of tuples $G_1(\vec{v}) \cup G_2(\vec{v})$ is consistent, as there are no tuples coinciding in the values of the attributes corresponding to $\vec{x}$, but not in the attributes corresponding to z_1 and z_2. It remains to be proved that, if these two tuples exist, then the tuples in $G_1(\vec{v}) \cup G_2(\vec{v})$ are consistent w.r.t. $\mathcal{IC}$. In fact, equation (A.7) ensures that $p(\tilde{t}_1) + p(\tilde{t}_2) \leq 1$, which in turn entails that a model for $\{\widetilde{t_1}, \widetilde{t_2}\}$ w.r.t. $\mathcal{IC}$ exists. Starting from this model, a model M for $G_1(\vec{v}) \cup G_2(\vec{v})$ w.r.t. $\mathcal{IC}$ can be obtained as follows. The tuples in $G_1(\vec{v})$ other than $\tilde{t}_1$ which are conflicting with at least one tuple $G_2(\vec{v})$ are put in the portion of the probability space occupied by the worlds containing $\widetilde{t_1}$. This can be done since the fact that $\tilde{t}_1$ has maximum probability among the tuples in $G_1(\vec{v})$ conflicting with at least one tuple in $G_2(\vec{v})$ makes any other tuple in $G_1(\vec{v})$ conflicting with at least one tuple in $G_2(\vec{v})$ have a probability which fits the portion of the probability space occupied by $\tilde{t}_1$. Similarly, the tuples in $G_2(\vec{v})$ other than $\tilde{t}_2$ which are conflicting with at least one tuple $G_1(\vec{v})$ are put in the portion of the probability space occupied by the worlds containing $\widetilde{t_2}$. Also in this case, this can be done since $\tilde{t}_2$ has maximum probability among the tuples in $G_2(\vec{v})$ conflicting with at least one tuple in $G_1(\vec{v})$. Finally, any tuple in $G_1(\vec{v})$ (resp., $G_2(\vec{v})$) which is conflicting with no tuple in $G_2(\vec{v})$ (resp., $G_1(\vec{v})$) can be put in any portion of the probability space, since its co-occurrence with any other tuple makes no constraint violated. □

Theorem 7. *There is an $\mathcal{IC}$ consisting of* 2 *FDs over the same relation scheme such that* CC *is NP-hard.*

Proof. We show a LOGSPACE reduction from 3-COLORING to CC which yields CC instances where $\mathcal{IC}$ contains only functional dependencies. The rationale of the proof is similar to the proof in [22] of the NP-hardness of PSAT.

We briefly recall the definition of 3-COLORING. An instance of 3-COLORING consists of a graph $G = \langle N, E\rangle$, where N is a set of node identifiers and E is a set of edges (pairs of node identifiers). The answer of a 3-COLORING instance is *true* iff there is a total function $f : N \to \{Red, Green, Blue\}$ such that $f(n_i) \neq f(n_j)$ whenever $\{n_i, n_j\} \in E$ (f is said to be a 3-coloring function over G).

Let $G = \langle N, E\rangle$ be a 3-COLORING instance. We construct an equivalent instance $\langle \mathcal{D}^p, \mathcal{IC}, D^p\rangle$ of CC as follows:

- $\mathcal{D}^p$ consists of the probabilistic relation schema $R^p(Node, Color, IdEdge, P)$;
- D^p is the instance of $\mathcal{D}^p$ consisting of the instance r^p of R^p defined as follows: for each node $n \in N$, for each edge $e \in E$ such that $n \in e$, and for each color $c \in \{Red,Green,Blue\}$, r^p contains the tuple $(n, c, e, \frac{1}{3})$.
- $\mathcal{IC}$ is the set of denial constraints over $\mathcal{D}^p$ consisting of the following two functional dependencies:

$ic_1 : \neg[R(x_1, x_2, x_3) \wedge R(x_1, x_4, x_5) \wedge x_2 \neq x_4]$
$ic_2 : \neg[R(x_1, x_2, x_3) \wedge R(x_4, x_2, x_3) \wedge x_1 \neq x_4]$

We first show that, if G is 3-colorable, then $D^p \models \mathcal{IC}$. In fact, given a 3-coloring function f over G, the interpretation Pr defined below is a model of D^p w.r.t. $\mathcal{IC}$. Pr assigns non zero probability to the following three possible worlds only:
$w_1 = \{R(n, f(n), e) \mid n \in N, e \in E \wedge n \in e\}$;
$w_2 = \{R(n, Next(f(n)), e) \mid n \in N, e \in E \wedge n \in e\}$;
$w_3 = \{R(n, Next(Next(f(n))), e) \mid n \in N, e \in E \wedge n \in e\}$,
where *Next* is a function which receives a color $c \in \{Red, Green, Blue\}$ and returns the next color in the sequence [*Red, Green, Blue*] (where *Next*(*Blue*) returns *Red*). Specifically, Pr assigns probability $\frac{1}{3}$ to all the three possible worlds w_1, w_2, w_3. It is easy to see that each possible world w_1, w_2, w_3 satisfies $\mathcal{IC}$ and that every tuple in D^p appears exactly in one possible world in $\{w_1, w_2, w_3\}$. Therefore Pr is a model of D^p.

We now show that, if $D^p \models \mathcal{IC}$, then G is 3-colorable. It is easy to see that G is 3-colorable if there is a model Pr for D^p w.r.t. $\mathcal{IC}$ having the following property Π: Pr assigns non-zero probability only to 3-*coloring possible worlds*, i.e., possible worlds containing, for each edge $e = (n_i, n_j) \in E$, two tuples $t_i^e = R(n_i, c_i, e)$ and $t_j^e = R(n_j, c_j, e)$, where $c_i \neq c_j$. In fact, starting from Pr and a 3-coloring possible world w with $Pr(w) > 0$, a function f^w can be defined which assigns to each node $n \in N$ the color c if there is a tuple $R(n, c, e) \in w$ (f^w is a function since it is injective, as w cannot contain tuples assigning different colors to the same node). Clearly, f^w is a 3-coloring function, as it associates every node n with a unique color and assigns different colors to pairs of nodes connected by an edge. Hence, it remains to be shown that at least one model satisfying Π exists. In fact, we prove that any model for D^p w.r.t. $\mathcal{IC}$ satisfies Π. Reasoning by contradiction, assume that, for a model Pr, there is a non-3-coloring possible world w^* such that $Pr(w^*) = \epsilon > 0$. That is, there is at least a pair n, e, with $n \in N$ and $e \in E$ such that for each $c \in \{Red,Green,Blue\}$, $R(n, c, e) \notin w^*$. Now, consider the tuples $t_1 = R(n,Red, e)$, $t_2 = R(n,Green, e)$, $t_3 = R(n,Blue, e)$ and the sets
$S_1 = \{w \in pwd(D^p) \mid t_1 \in w \wedge Pr(w) > 0\}$,
$S_2 = \{w \in pwd(D^p) \mid t_2 \in w \wedge Pr(w) > 0\}$,
$S_3 = \{w \in pwd(D^p) \mid t_3 \in w \wedge Pr(w) > 0\}$.
Since ic_1 is satisfied by every possible world $w \in pwd(D^p)$ such that $Pr(w) > 0$, this means that for each possible world w there is at most one color $c \in \{Red,Green,Blue\}$ such that the tuple $R(n, c, e)$ belongs to w. Therefore, it must be the case that, $\forall i, j \in \{1, 2, 3\}, i \neq j$, $S_i \cap S_j = \emptyset$. Since Pr is an interpretation, the following equalities must hold:

- $\frac{1}{3} = p(t_1) = \sum_{w \in S_1} Pr(w)$;
- $\frac{1}{3} = p(t_2) = \sum_{w \in S_2} Pr(w)$;
- $\frac{1}{3} = p(t_3) = \sum_{w \in S_3} Pr(w)$.

This implies that

$$\sum_{w \in S_1} Pr(w) + \sum_{w \in S_2} Pr(w) + \sum_{w \in S_3} Pr(w) = 1$$

However, since $Pr(w^*) = \epsilon > 0$ and Pr is an interpretation, $\sum_{w \in pwd(D^p) \setminus \{w^*\}} Pr(w) < 1$. The latter, since $w^* \notin S_i$ for each $i \in \{1, 2, 3\}$, implies that $pwd(D^p) \setminus \{w^*\} \supseteq S_1 \cup S_2 \cup S_3$, and then $\sum_{w \in (S_1 \cup S_2 \cup S_3)} Pr(w) < 1$ which is a contradiction. □

Theorem 8. *Let each denial constraint in $\mathcal{IC}$ be join-free or a BEGD. If, for each pair of distinct constraints ic_1,ic_2 in $\mathcal{IC}$, the relation names occurring in ic_1 are distinct from those in ic_2, then* CC *is in PTIME.*

Proof. Trivially follows from theorems 6, 4, and from the fact that the consistency can be checked by considering the maximal connected components of the conflict hypergraph separately. □

Theorem 9. *If $\mathcal{IC}$ consists of one FD per relation, then $HG(D^p, \mathcal{IC})$ is a graph where each connected component is either a singleton or a complete multipartite graph. Moreover, D^p is consistent w.r.t. $\mathcal{IC}$ iff the following property holds: for each connected component C of $HG(D^p, \mathcal{IC})$, denoting the maximal independent sets of C as $S_1, \ldots, S_k$, it is the case that $\sum_{i \in [1..k]} \tilde{p}_i \leq 1$, where $\tilde{p}_i = \max_{t \in S_i} p(t)$.*

Proof. It is easy to see that multiple FDs over distinct relations involve disjoint sets of tuples. Thus, it is straightforward to see that the conflict hypergraph has the structural property described in the statement iff, for each relation, the conflict hypergraph over the set of tuples of this relation is a graph having the same structural property. Moreover, as observed in the proof of Theorem 8, the consistency can be checked by considering the maximal connected components of the conflict hypergraph separately.

This implies that, in order to prove the statement, it suffices to consider the case that that $\mathcal{IC}$ consists of a unique FD ic over a relation R, and D^p consists of an instance r of R. In particular, we assume that ic is of the form:

$$\neg[R(\vec{x}, \vec{y}_1) \wedge R(\vec{x}, \vec{y}_2) \wedge z_1 \neq z_2],$$

where z_1 and z_2 are variables in $\vec{y}_1$ and $\vec{y}_2$, respectively, corresponding to the same attribute Z of R. That is, we are assuming that the FD ic is in canonical form (i.e., its right-hand side consists of a unique attribute). This yields no loss of generality, as it is easy to see that the reasoning used in the proof is still valid in the presence of FDs whose right-hand sides contain more than one attribute.

The relation instance r can be partitioned into the two relations r', r'', containing the tuples connected to at least another tuple in $HG(D^p, \mathcal{IC})$ (that is, tuples belonging to some conflicting set) and the isolated tuples (that is, tuples belonging to no conflicting set), respectively. Obviously, the subgraph of $HG(D^p, \mathcal{IC})$ containing only the tuples in r'' contains no edge, and it is such that each of its connected component is a singleton. Therefore, in order to complete the proof of the first part of the statement, it remains to be proved that the subgraph G of $HG(D^p, \mathcal{IC})$ containing only the tuples in r' is such that each of its connected component is a complete multipartite graph.

Let $\overline{X}$ be the set of attributes in $Attr(R)$ corresponding to the variables in $\vec{x}$. The form of ic implies that G is a graph having the following structural property $\mathcal{S}$: for any pair of tuples t_1, t_2, there is the edge (t_1, t_2) in G iff: 1) $\forall X \in \overline{X}$, $t_1[X] = t_2[X]$, and 2) $t_1[Z] \neq t_2[Z]$.

This implies that G has as many connected components as the cardinality of $\Pi_{\overline{X}} r'$. Specifically, each connected component of G corresponds to a tuple $\vec{v}$ in $\Pi_{\overline{X}} r'$, as it contains every tuple of r' whose projection over $\overline{X}$ coincides with $\vec{v}$. In fact, property $\mathcal{S}$ implies that:

A. there is no path in G between tuples differing in at least one attribute in $\overline{X}$;
B. any two tuples t', t'' coinciding in all the attributes in $\overline{X}$ are either directly connected to one another (in the case that they do not coincide in attribute Z), or there is a third tuple t''' to which they are both connected. In fact, t' and t'' are not isolated (otherwise they would not belong to r'), and any tuple conflicting with t' is also conflicting with t'', as we are in the case that t' and t'' coincide in Z.

To complete the proof of the first part of the statement, we now show that, taken any connected component C of G, C is a complete multipartite graph. This straightforwardly follows from the following facts:

a. the nodes of C can be partitioned into the maximal independent sets $S_1, \ldots, S_k$, where k is the number of distinct values of attribute Z occurring in the tuples in C. In particular, each S_i corresponds to one of these values v of Z, and contains all the tuples of C having v as value of attribute Z. The fact that every S_i is a maximal independent set trivially follows from property $\mathcal{S}$.

b. for every pair of tuples t_i and t_j belonging to S_i and S_j (with $i, j \in [1..k]$ and $i \neq j$), there is an edge connecting t_i to t_j (this also trivially follows from property $\mathcal{S}$).

We now prove the second part of the statement.

($\Rightarrow$): Reasoning by contradiction, assume that D^p is consistent w.r.t. $\mathcal{IC}$ but, for some connected component C of $HG(D^p, \mathcal{IC})$, it does not hold that $\sum_{i\in[1..k]} \tilde{p}_i \leq 1$, where $\tilde{p}_i = \max_{t\in S_i} p(t)$ and $S_1, \ldots, S_k$ are the maximal independent sets of C. Obviously, C can not be a singleton (otherwise the inequality would hold), thus it must be the case that C is a complete multipartite graph.

For each $i \in [1..k]$, let $\tilde{t}_i$ be a tuple of S_i such that $p(\tilde{t}_i) = \tilde{p}_i$. Since C is a complete multipartite graph, and since the so obtained tuples $\tilde{t}_1, \ldots, \tilde{t}_k$ belong to distinct independent sets, it must be the case that, for each $i, j \in [1..k]$ with $i \neq j$, there is an edge in C between $\tilde{t}_i$ and $\tilde{t}_j$. This means that, in every model M for D^p w.r.t. $\mathcal{IC}$, for each $i, j \in [1..k]$ with $i \neq j$, the tuples $\tilde{t}_i, \tilde{t}_j$ can not co-exist in a non-zero probability possible world. That is, every non-zero probability possible world contains at most one tuple among those in $\{\tilde{t}_1, \ldots, \tilde{t}_k\}$. This entails that the sum of the probabilities of the possible worlds containing the tuples in $\{\tilde{t}_1, \ldots, \tilde{t}_k\}$ is equal to the sum of the marginal probabilities of the tuples in $\{\tilde{t}_1, \ldots, \tilde{t}_k\}$, which, by contradiction hypothesis, is greater than 1. This contradicts the fact that M is a model.

($\Leftarrow$): We now show that D^p is consistent w.r.t. $\mathcal{IC}$ if the inequality $\sum_{i\in[1..k]} \tilde{p}_i \leq 1$ holds, where $\tilde{p}_i = \max_{t\in S_i} p(t)$ and $S_1, \ldots, S_k$ are the maximal independent sets of C. Consider the database instance $\tilde{D}^p$ consisting of the tuples $\tilde{t}_1, \ldots, \tilde{t}_k$ where $\tilde{t}_i$ (with $i \in [1..k]$) is a tuple of S_i such that $p(\tilde{t}_i) = \tilde{p}_i$. It is easy to see that there is a model for $\tilde{D}^p$ w.r.t. IC: since C is a complete multipartite graph, and $\tilde{t}_1, \ldots, \tilde{t}_k$ belong to distinct independent sets of C, it follows that, for each $i, j \in [1..k]$ with $i \neq j$, there is exactly one edge in C between $\tilde{t}_i$ and $\tilde{t}_j$. That is, the conflict graph of $\tilde{D}^p$ w.r.t. $\mathcal{IC}$ is a clique. Hence, the fact that inequality $\sum_{i\in[1..k]} \tilde{p}_i \leq 1$ holds is sufficient to ensure the existence of a model $\tilde{M}$ for $\tilde{D}^p$ w.r.t. $\mathcal{IC}$. Starting from $\tilde{M}$, a model for D^p w.r.t. $\mathcal{IC}$ can be obtained by reasoning as follows. Since, for each maximal independent set S_i of C (with $i \in [1..k]$), the tuples in S_i other than $\tilde{t}_i$ are such that their probability is less than or equal to $p(\tilde{t}_i)$, a model M for D^p w.r.t. $\mathcal{IC}$ can be obtained by putting the tuples in S_i other than $\tilde{t}_i$ in the portion of the probability space corresponding to that occupied by the worlds containing $\tilde{t}_i$ according the model $\tilde{M}$. The fact that M is a model follows from the fact that, for each $i \in [1..k]$, the tuples in S_i other than $\tilde{t}_i$ are conflicting only with the same tuples which are conflicting with $\tilde{t}_i$. □

Appendix A.4. Proofs of Lemma 2 and Theorem 11

Lemma 2. *Let Q be a conjunctive query over $\mathcal{D}^p$, D^p an instance of $\mathcal{D}^p$, and $\vec{t}$ an answer of Q having minimum probability p^{min} and maximum probability p^{max}. Let m be the number of tuples in D^p plus 3 and a be the maximum among the numerators and denominators of the probabilities of the tuples in D^p. Then p^{min} and p^{max} are expressible as fractions of the form $\frac{\eta}{\delta}$, with $0 \leq \eta \leq (ma)^m$ and $0 < \delta \leq (ma)^m$.*

Proof. Consider the equivalent form of the linear programming problem $LP(S^*)$ described in the proof of Proposition 2, where equalities (e1) of $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ are rewritten as:
$\forall t \in D^p,\ d(p(t)) \times \sum_{i|w_i\in pwd(D^p)\wedge t\in w_i} v_i = d(p(t)) \times p(t)$,
where $p(t) = \frac{n(p(t))}{d(p(t))}$ (i.e., $n(p(t))$ and $d(p(t))$ are the numerator and denominator of $p(t)$, respectively). This way, we have that all the coefficients of $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ are integers, where each coefficient can be either 0, or 1, or the numerator or the denominator of the marginal probability of a tuple of D^p.

In [42], it was shown that the solution of any instance of the linear programming problem with integer coefficients is expressible as a fraction of the form $\frac{\eta}{\delta}$, where both η and δ are naturals bounded by $(ma)^m$, where m is the number of (in)equalities and a the greatest integer coefficient occurring in the instance. By applying this result to $LP(S^*)$, we get the statement: in fact, it is easy to see that *i*) $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ contains integer coefficients only, *ii*) the number m of (in)equalities in $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ is equal to the number of tuples in D^p plus 3, and *iii*) the greatest integer constant a in $S^*(\mathcal{D}^p, \mathcal{IC}, D^p)$ is the maximum among the numerators and denominators of the probabilities of the tuples in D^p. □

Theorem 11. (Lower bound of MP) *There is at least one conjunctive query without projection for which* MP *is coNP-hard, even if $\mathcal{IC}$ consists of binary constraints only.*

Proof. We show a reduction from the planar 3-COLORING problem to the complement of the membership problem ($\overline{\text{MP}}$). An instance of planar 3-COLORING consists of a planar graph $G = \langle N, E\rangle$, where N is a set of node identifiers and E is a set of edges (pairs of node identifiers). The answer of a planar 3-COLORING instance G is true iff there is a 3-coloring function over G, i.e., a total function $f : N \to \{R, G, B\}$ such that $f(n_i) \neq f(n_j)$ whenever $\{n_i, n_j\} \in E$. Observe that every planar graph $G = \langle N, E\rangle$ is 4-colorable, that is, there exists a function $f : N \to \{R, G, B, C\}$ such that $f(n_i) \neq f(n_j)$ whenever $\{n_i, n_j\} \in E$ (in this case, f is said to be a 4-coloring function).

Let $G = \langle N, E\rangle$ be a planar 3-COLORING instance. We construct an equivalent $\overline{\text{MP}}$ instance $\langle \mathcal{D}^p, \mathcal{IC}, D^p, Q, t, k_1, k_2\rangle$ as follows:

- $\mathcal{D}^p$ consists of the probabilistic relation schemas R^p_G(*Node, Color, IdEdge, P*) and R^p_ϕ(*Tid,P*).
- D^p is the instance of $\mathcal{D}^p$ consisting of the instances r^p_G of R^p_G and r^p_ϕ of R^p_ϕ defined as follows:
 - for each node $n \in N$ and for each edge $e \in E$ such that $n \in e$, r^p_G contains four tuples of the form $R^p_G(n, c, e, \frac{1}{8})$, one for each $c \in \{R, G, B, C\}$;
 - r^p_ϕ consists of the tuples $R^p_\phi(1, \frac{1}{2})$ and $R^p_\phi(2, \frac{1}{2})$ only;
- $\mathcal{IC}$ contains the following binary denial constraints:

$ic_1 : \neg[R_G(x_1, x_2, x_3) \wedge R_G(x_1, x_4, x_5) \wedge x_2 \neq x_4]$;
$ic_2 : \neg[R_G(x_1, x_2, x_3) \wedge R_G(x_4, x_2, x_3) \wedge x_1 \neq x_4]$;
$ic_3 : \neg[R_G(x_1, x_2, x_3) \wedge R_\phi(2)]$;
$ic_4 : \neg[R_G(x_1, x_2, C) \wedge R_\phi(1)]$;

- $Q(x, y) = R_\phi(x) \wedge R_\phi(y)$;
- $t = (1, 2)$;
- $k_1 = \frac{1}{2}$;
- $k_2 = 1$.

It is easy to see that the fact that G is 4-colorable implies that D^p is consistent w.r.t. $\mathcal{IC}$ (it suffices to follow the same reasoning as the proof of Theorem 1, using 4 colors instead of 3).

We first prove that, if G is 3-colorable, then the corresponding instance of $\overline{\text{MP}}$ is true. Let f be a 3-coloring function over G. Consider an interpretation Pr for D^p which assigns non-zero probability to the following possible worlds only:
$w_1 = \{R_G(n, f(n), e) \mid n \in N, e \in E \wedge n \in e\} \cup \{R_\phi(1)\}$
$w_2 = \{R_G(n, \mathit{Next}(f(n)), e) \mid n \in N, e \in E \wedge n \in e\}$
$w_3 = \{R_G(n, \mathit{Next}(\mathit{Next}(f(n))), e) \mid n \in N, e \in E \wedge n \in e\}$
$w_4 = \{R_G(n, \mathit{Next}(\mathit{Next}(\mathit{Next}(f(n)))), e) \mid n \in N, e \in E \wedge n \in e\}$
$w_5 = \{R_\phi(1), R_\phi(2)\}$
$w_6 = \{R_\phi(2)\}$
where *Next* is a function which receives a color $c \in \{R, G, B, C\}$ and returns the next color in the sequence $[R, G, B, C]$ (where $\mathit{Next}(C)$ returns R). Furthermore, Pr assigns probability $\frac{1}{8}$ to the possible worlds w_1, w_2, w_3, w_4 and w_5, and probability $\frac{3}{8}$ to the possible world w_6. It is easy to see that Pr is a model of D^p w.r.t. $\mathcal{IC}$ and the probability that the tuple $t = (1, 2)$ is an answer of Q assigned by Pr is $\frac{1}{8}$. Hence, the $\overline{\text{MP}}$ is true in this case (as $\frac{1}{8} < k_1$).

We now prove that if G is not 3-colorable, then the corresponding instance of $\overline{\text{MP}}$ is false. First observe that, reasoning similarly to in the proof of Theorem 1, it is possible to show that, for each model Pr of D^p w.r.t. $\mathcal{IC}$ and for each possible world w such that $Pr(w) > 0$, if w contains at least a tuple of r_G, then for each node $n \in N$ and for each edge $e \in E$ such that $n \in e$, there exists $c \in \{R, G, B, C\}$ such that w contains the tuple $R_G(n, c, e)$. This is due to the fact that every possible world w such that $Pr(w) > 0$ can not contain two tuples $R_G(n, c', e), R_G(n, c'', e)$ and no tuple in r_G can belong to a possible world which contains the tuple $R_\phi(2)$ too.

Since G is not 3-colorable, for each model Pr of D^p w.r.t. $\mathcal{IC}$ and for each possible world w such that $Pr(w) > 0$ containing at least a tuple of r_G, it holds that w contains a tuple $R_G(n, C, e)$. This implies that no possible world containing a tuple of r_G can contain the tuple $R_\phi(1)$, as otherwise ic_4 would be violated. Since ic_1 and ic_3 hold for Pr, then the sum of the probability of the possible worlds containing at least a tuple of r_G is equal to $\frac{1}{2}$. Since the possible worlds containing at least a tuple of r_G cannot contain neither $R_\phi(1)$ nor $R_\phi(2)$ (as ic_4 holds) and both $R_\phi(1)$ nor $R_\phi(2)$ has probability $\frac{1}{2}$ it holds that the probability that both $R_\phi(1)$ and $R_\phi(2)$ are true is $\frac{1}{2}$. The latter implies that the minimum probability that $t = (1, 2)$ is an answer of Q is $\frac{1}{2}$, which is equal to k_1. Therefore the $\overline{\text{MP}}$ is false if G is not 3-colorable. □

Appendix A.5. Proof of Theorem 12

Theorem 12. (QA complexity) QA *belongs to* FP^{NP} *and is* $FP^{NP[\log n]}$*-hard.*

Proof. The membership in FP^{NP} follows from [35], where it was shown that a problem more general than ours (that is, the entailment problem for probabilistic logic programs with conditional rules) belongs to FP^{NP} (see *Related Work*). We prove the hardness for $FP^{NP[\log n]}$ by showing a reduction to QA from the well-known $FP^{NP[\log n]}$-hard problem CLIQUE SIZE, that is the problem of determining the size K^* of the largest clique of a given graph.

Let the graph $G = \langle N, E\rangle$ be an instance of CLIQUE SIZE, where $u_1, \ldots, u_n$ are the nodes of G (where $n = |N|$). We construct an equivalent instance $\langle \mathcal{D}^p, \mathcal{IC}, D^p, Q\rangle$ of QA as follows. D^p is the database schema consisting of the following relation schemas: $Node^p(Id, P)$, $NoEdge^p(nodeId_1, nodeId_2, P)$, $Flag^p(Id, P)$. The database instance D^p consists of the following relation instances. Relation $node^p$ contains a tuple $t_i = Node^p(u_i, 1/n)$ for each node u_i of G (that is, every node of G corresponds to a tuple of $node^p$ having probability $1/n$). Relation $noEdge^p$ contains a tuple $NoEdge^p(u_i, u_j, 1)$ for each pair of distinct nodes of G which are not connected by means of any edge in E (thus, $noEdge^p$ represents the complement of E, and all of its tuples have probability 1). Finally, relation $flag^p$ contains the unique tuple $Flag^p(1, \frac{n-1}{n})$.

Let $\mathcal{IC}$ consist of the following denial constraints over $\mathcal{D}^p$:

$$ic_1 : \neg[Node(x_1) \wedge Node(x_2) \wedge NoEdge(x_1, x_2)]$$

$$ic_2 : \neg[Node(x_1) \wedge Node(x_2) \wedge Flag(1) \ \wedge \ x_1 \neq x_2]$$

Basically, constraint ic_1 forbids that tuples representing distinct nodes co-exist if they are not connected by any edge, while ic_2 imposes that tuple $Flag(1)$ can co-exist with at most one tuple representing a node.

To complete the definition of the instance of QA, we define the (boolean) query $Q() = Flag(1) \wedge Node(x)$.

We will show that the size of the largest clique of G is K^* iff the empty tuple $t_\emptyset$ is an answer of Q over D^p with minimum probability $l^* = \frac{n-K^*}{n}$ (i.e., $Ans(Q, D^p, \mathcal{IC})$ consists of the pair $\langle t_\emptyset, [p^{min}, p^{max}]\rangle$, with $p^{min} = \frac{n-K^*}{n}$).

We first show that if G contains a clique of size K, then $p^{min} \leq \frac{n-K}{n}$. In fact, if K is the size of a clique C of G, then we can construct the following model M for D^p w.r.t. $\mathcal{IC}$. Let $w^c = \{Node(u_i) | u_i \in C\} \cup noEdge$, $w^f = \{ Flag(1)\} \cup noEdge$, and, for each $u_i \in N \setminus C$, $w_i = \{Node(u_i), Flag(1)\} \cup noEdge$. Then, denoting as w the generic possible world, the model M is defined as follows:

$$M(w) = \begin{cases} 1/n & \text{if } w = w^c; \\ 1/n & \text{if } w = w_i, \text{ for } i \text{ s.t. } u_i \in N \setminus C; \\ (K-1)/n & \text{if } w = w^f; \\ 0 & \text{otherwise} \end{cases}$$

It is easy to see that M is a model. First of all, it assigns non-zero probability only to possible worlds satisfying the constraints. Moreover, for any tuple t in D^p, summing the probabilities of the possible worlds containing t results in $t[P]$. In fact, considering only the possible worlds which have been assigned a non-zero probability by M, every tuple $Node(u_i)$ representing a node $u_i \in C$ belongs only to w^c, which is assigned by M the probability $1/n$ (the same as $p(Node(u_i))$). Analogously, every tuple $Node(u_i)$ representing a node $u_i \notin C$ belongs only to w^i, which is assigned by M the probability $1/n$ (the same as $p(Node(u_i))$). Finally, tuple $Flag(1)$ occurs only in w^f and in $n - K$ possible worlds of the form w_i, thus the sum of the probabilities of the possible worlds containing $Flag(1)$ is $M(w^f) + (n - K) \cdot \frac{1}{n} = \frac{(n-1)}{n} = p(Flag(1))$.

It is easy to see that the probability of the answer $t_\emptyset$ of Q over the model M is the sum of the probabilities of the possible worlds of the form w_i, that is $\frac{(n-K)}{n}$. Hence, from definition of minimum probability, it holds that $p^{min} \leq \frac{(n-K)}{n}$.

To complete the proof, it suffices to show that the following property $\mathcal{P}$ holds over any model M' for $\mathcal{D}^p$ w.r.t. $\mathcal{IC}$: "*the probability l of the answer $t_\emptyset$ of Q over M' can not be strictly less than $l^* = \frac{(n-K^*)}{n}$*". Observe that, for every model M', the possible worlds which have been assigned a non-zero probability by M' can be of three types (we do not consider *noEdge* tuples, as they have probability 1, thus they belong to every non-zero-probability possible world):

Type 1: world not containing *Flag*(1), and containing a non-empty set of tuples representing the nodes of a clique (the non-emptiness of this set derives from the combination of constraint ic_2 with the value of the marginal probability assigned to tuple *Flag*(1));

Type 2: world containing the tuple *Flag*(1) and exactly one *node* tuple;

Type 3: world containing the tuple *Flag*(1) only.

We will show that property $\mathcal{P}$ holds over any model M' by reasoning inductively on the number x of possible worlds of Type 1 which have been assigned a non-zero probability by M'.

The base case is $x = 1$, meaning that M' assigns probability $1/n$ to a unique Type-1 world w_1^{T1}, and probability 0 to all the other possible worlds of the same type. It is easy to see that the sum of the probabilities assigned by M' to the Type-2 worlds (which coincides with l) is equal to $\frac{1}{n} \cdot (n - |C_1^{T1}|)$, where C_1^{T1} is the clique represented by w_1^{T1}. Hence, if it were $l < l^*$, it would hold that $\frac{1}{n} \cdot (n - |C_1^{T1}|) < \frac{(n-K^*)}{n}$, which means that $|C_1^{T1}| > K^*$, thus contradicting that K^* is the size of the maximum clique of G.

We now prove the induction step. The induction hypothesis is that $\mathcal{P}$ holds over any model assigning non-zero probability to exactly $x-1$ Type-1 possible worlds (with $x - 1 \geq 1$). We show that this implies that $\mathcal{P}$ holds also over any model assigning non-zero probability to exactly x Type-1 possible worlds. Consider a model M' assigning non-zero probability to exactly x Type-1 possible worlds, namely $w_1^{T1}, \ldots, w_x^{T1}$. We assume that these worlds are ordered by their cardinality (in descending order), and denote as C_i the clique represented by w_i^{T1} (with $i \in [1..x]$). We also denote as $w_1^{T2}, \ldots, w_n^{T2}$ the Type-2 possible worlds (where w_i^{T2} contains the *node* tuple representing u_i). Moreover, let l' be the probability of the answer $t_\emptyset$ of Q over M'. We show that, starting from M', a new model M'' for D^p w.r.t. $\mathcal{IC}$ can be constructed such that:
i) M'' assigns non-zero probability to $x - 1$ Type-1 possible worlds;
ii) the probability l'' of the answer *true* of Q over M'' satisfies $l'' \leq l'$.
Specifically, M'' is defined as follows. M'' coincides with M' on all the Type-1 worlds except for the probabilities assigned to w_1^{T1} and w_x^{T1}. In particular, $M''(w_1^{T1}) = M'(w_1^{T1}) + M'(w_x^{T1})$, while $M''(w_x^{T1}) = 0$. Moreover, for each Type-2 world w_i^{T2} such that $u_i \in C_1 \setminus C_x$, $M''(w_i^{T2}) = M'(w_i^{T2}) - M'(w_x^{T1})$, and, for each Type-2 world w_i^{T2} such that $u_i \in C_x \setminus C_1$, $M''(w^{T2}) = M'(w^{T2}) + M'(w_x^{T1})$. On the remaining Type-2 worlds, M'' is set equal to M'. Finally, denoting the type-3 world as w^{T3}, $M''(w^{T3}) = M'(w^{T3}) - |C_x \setminus C_1| \cdot M'(C_x) + |C_1 \setminus C_x| \cdot M'(C_x)$. In brief, M'' is obtained from M' by moving the probability assigned to w_x^{T1} to w_1^{T1}, and re-assigning the probabilities of the Type-2 and Type-3 worlds accordingly. Hence, it is easy to see that M'' is still a model (as it can be easily checked that it makes the sum of the probabilities of the possible worlds containing a tuple equal to the marginal probability of the tuple). Moreover, property *i*) holds, as M'' assigns probability 0 to the world w_x^{T1}, while the other worlds of the form w_i^{T1} (with $i < x$) are still assigned by M'' a positive probability, and the remaining Type-1 worlds are still assigned probability 0. Also property *ii*) holds, since the probability of *true* as answer of Q over M'' is given by $l'' = l' + |C_x \setminus C_1| \cdot M'(C_x) - |C_1 \setminus C_x| \cdot M'(C_x)$. Since $|C_1| \geq |C_x|$, and thus $|C_1 \setminus C_x| \geq |C_x \setminus C_1|$, l'' is less than or equal to l'. If it were $l' < l^*$ (and thus $l'' < l^*$) M'' would be a model assigning non-zero probability to $x-1$ Type-1 possible worlds such that the answer *true* of Q over M'' has probability strictly less than l^*, thus contradicting the induction hypothesis. □

Appendix A.6. Proof of Theorem 13

The proof of Theorem 13 is postponed to the end of this section, after introducing some preliminary lemmas.

Lemma 4. *Let D^p be a PDB instance of $\mathcal{D}^p$ such that $HG(D^p, \mathcal{IC})$ is a graph and $D^p \models \mathcal{IC}$. Let t, t' be two tuples connected by exactly one path in $HG(D^p, \mathcal{IC})$. Then, $p^{min}(t \wedge t')$ and $p^{max}(t \wedge t')$ can be computed in polynomial time w.r.t. the size of D^p.*

Proof. Let π be the path connecting t and t' in $HG(D^p, \mathcal{IC})$. It is easy to see that the fact that π is unique implies that $p^{min}(t \wedge t') = p_\pi^{min}(t \wedge t')$ and $p^{max}(t \wedge t') = p_\pi^{max}(t \wedge t')$ (in fact, any model for D^p w.r.t. $HG(D^p, \mathcal{IC})$ can be obtained by refining a model for D^p w.r.t. π without changing the probabilities assigned to the event $t \wedge t'$, following a reasoning analogous to that used in the proof of the right-to-left implication of Theorem 2).

Since the path π connecting t and t' in the graph $HG(D^p, \mathcal{IC})$ is unique, it does not contain cycles (otherwise there would be at least two paths between t and t'). Hence, π is a chain in a graph (the definition of chain for hypergraph is introduced in Section Appendix A.2). Therefore, $p_\pi^{min}(t \wedge t')$ can be determined by exploiting Lemma 3, which provides the formula for computing the minimum probability that the ears at the endpoints of a chain co-exist. It is trivial to see that, denoting as $\hat{t}$ and $\hat{t}'$ the tuples connected to t and t' in π, in our case the formula in Lemma 3 becomes:

$$p_\pi^{min}(t \wedge t') = \begin{cases} 0, \text{ if } (t,t') \text{ is an edge of } \pi \\ \max\{0, p(t)+p(t')-[1-p_\pi^{min}(\hat{t} \wedge \hat{t}')]\}, \text{ otherwise.} \end{cases}$$

since π is a chain in a graph, thus its intermediate edges are hyperedges of cardinality 2 with no ears.

As regards $p_\pi^{max}(t \wedge t')$, it can be evaluated as follows:

$$p_\pi^{max}(t \wedge t') = \begin{cases} 0, \text{ if } (t,t') \text{ is an edge of } \pi \\ \min\{p(t), p(t'), 1-[p(\hat{t})+p(\hat{t}')-p_\pi^{max}(\hat{t} \wedge \hat{t}')]\}, \\ \quad \text{otherwise.} \end{cases}$$

In fact, it is easy to see that the maximum probability of the event $t \wedge t'$ is $\min\{p(t), p(t'), p_\pi^{max}(\neg\hat{t} \wedge \neg\hat{t}')\}$, where $p_\pi^{max}(\neg\hat{t} \wedge \neg\hat{t}')$ is the maximum probability that both the tuples $\hat{t}$ and $\hat{t}'$ (which are mutually exclusive with t and t', respectively) are false. In turn, $p_\pi^{max}(\neg\hat{t} \wedge \neg\hat{t}') = 1 - p_\pi^{min}(\hat{t} \vee \hat{t}') = 1 - [p(\hat{t}) + p(\hat{t}') - p_\pi^{max}(\hat{t} \wedge \hat{t}')]$, thus proving the above-reported formula.

We complete the proof by observing that $p^{min}(t \wedge t')$ and $p^{max}(t \wedge t')$ can be computed in polynomial time w.r.t. the size of D^p by recursively applying the above-reported formulas for p^{min} and p^{max} starting from t and t', and going further on towards the center of the unique path connecting t and t'. □

Lemma 5. *For projection-free queries,* QA *is in PTIME if $HG(D^p, \mathcal{IC})$ is a clique.*

Proof. It straightforwardly follows from the fact that, for each pair of tuples t, t' in $HG(D^p, \mathcal{IC})$, it holds that $p^{min}(t \wedge t') = p^{max}(t \wedge t') = 0$. □

Lemma 6. *For projection-free queries,* QA *is in PTIME if $HG(D^p, \mathcal{IC})$ is a tree.*

Proof. $Ans(Q, D^p, \mathcal{IC})$ can be determined by first evaluating the answer r_q of Q w.r.t. $det(D^p)$, and then computing, for each $\vec{t} \in r_q$, the minimum and maximum probabilities p^{min} and p^{max} of $\vec{t}$ as answer of Q. Obviously, r_q can be evaluated in polynomial time w.r.t. the size of D^p, and the number of tuples in r_q is polynomially bounded by the size of D^p.

Observe that, every ground tuple $\vec{t} \in r_q$ derives from the conjunction of a set of tuples $\{t_1, \ldots, t_n\}$ in $det(D^p)$. Thus, in order to prove the statement, it suffices to prove that, for each set $\{t_1, \ldots, t_n\}$ of tuples in $det(D^p)$, computing $p^{min}(t_1 \wedge \cdots \wedge t_n)$ and $p^{max}(t_1 \wedge \cdots \wedge t_n)$ is feasible in polynomial time w.r.t. the size of D^p.

For the sake of clarity of presentation, we assume that $HG(D^p, \mathcal{IC})$ coincides with its own minimal spanning tree containing all the tuples in $\{t_1, \ldots, t_n\}$. This means that each t_i (with $i \in [1..n]$) is either a leaf node or occurs as intermediate node in the path connecting two other tuples in $\{t_1, \ldots, t_n\}$, and all the leaf nodes are in $\{t_1, \ldots, t_n\}$. In fact, if this were not the case, it is straightforward to see that nothing would change in evaluating $p^{min}(t_1 \wedge \cdots \wedge t_n)$ and $p^{max}(t_1 \wedge \cdots \wedge t_n)$ if we disregarded the nodes of $HG(D^p, \mathcal{IC})$ which are not in $\{t_1, \ldots, t_n\}$ and do not belong to any path connecting some pair of nodes in $\{t_1, \ldots, t_n\}$.

Before showing how $p^{min}(t_1 \wedge \cdots \wedge t_n)$ and $p^{max}(t_1 \wedge \cdots \wedge t_n)$ can be computed, we introduce some notations. We say that a tuple t is a *branching node* of $HG(D^p, \mathcal{IC})$ iff the degree of t is greater than two. Moreover, a pair of tuples (t_i, t_j) is said to be an *elementary pair of tuples* of $HG(D^p, \mathcal{IC})$ if (*i*) each of t_i and t_j is either in $\{t_1, \ldots, t_n\}$ or a branching node, and (*ii*) the path connecting t_i to t_j contains neither branching nodes nor tuples in $\{t_1, \ldots, t_n\}$ as intermediate nodes.

The set of the elementary pairs of tuples is denoted as $EP_{HG(D^p, \mathcal{IC})}$ (we also use the short notation EP, when $HG(D^p, \mathcal{IC})$ is understood). Moreover, we denote the branching nodes of $HG(D^p, \mathcal{IC})$ which are not in $\{t_1, \ldots, t_n\}$ as $t_{n+1}, \cdots, t_{n+m}$. Observe that $m < n$, as n is also greater than or equal to the number of leaves of $HG(D^p, \mathcal{IC})$. Finally,

we denote with $B = \{true, false\}$ the boolean domain, with B^{n+m} the set of all the tuples of $n+m$ boolean values, and use the symbol α for tuples of $n+m$ boolean values and the notation $\alpha[i]$ to indicate the value of the i-th attribute of α.

We will show that $p^{min}(t_1 \wedge \cdots \wedge t_n)$ (resp., $p^{max}(t_1 \wedge \cdots \wedge t_n)$) is a solution of the following linear programming problem instance $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$:

$$\begin{array}{ll} \textit{minimize (resp., maximize)} & \displaystyle\sum_{\alpha \in B^{n+m} \mid \forall i \in [1..n]\, \alpha[i]=true} x_\alpha \\ \textit{subject to} & S(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p) \end{array}$$

where $S(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ is the following system of linear inequalities:

$$\begin{cases} \forall (t_i, t_j) \in EP & \displaystyle\sum_{\substack{\alpha \in B^{n+m} \mid \\ \alpha[i] = true \wedge \\ \alpha[j] = true}} x_\alpha = x_{t_i,t_j} & (A) \\ \forall (t_i, t_j) \in EP & p^{min}(t_i \wedge t_j) \leq x_{t_i,t_j} \leq p^{max}(t_i \wedge t_j) & (B) \\ \forall i \in [1..n+m] & \displaystyle\sum_{\alpha \in B^{n+m} \mid \alpha[i]=true} x_\alpha = p(t_i) & (C) \\ & \displaystyle\sum_{\alpha \in B^{n+m}} x_\alpha = 1 & (D) \end{cases}$$

Therein: (*i*) x_{t_i,t_j} is a variable representing the probability that t_i and t_j coexist; and (*ii*) $\forall \alpha \in B^{n+m}$, x_α is a variable representing the probability that $\forall i \in [1..n+m]$ the truth value of t_i is $\alpha[i]$; that is, x_α is the probability of the event $\bigwedge_{i \mid \alpha[i]=true} t_i \wedge \bigwedge_{i \mid \alpha[i]=false} \neg t_i$.

Since $HG(D^p, \mathcal{IC})$ is a tree, Lemma 4 ensures that, for each $(t_i, t_j) \in EP$, $p^{min}(t_i \wedge t_j)$ and $p^{max}(t_i \wedge t_j)$ can be computed in polynomial time w.r.t. the size of D^p. Therefore, we assume that they are precomputed constants in $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$.

It is easy to see that $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ can be solved in polynomial time w.r.t. the size of D^p, as it consists of at most $6n-2$ inequalities using $2^{2n-1} + 2n - 1$ variables, and n only depends on the number of relations appearing in Q (we recall that we are addressing data complexity, thus queries are of constant arity).

We now show that, for each solution of $S(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$, there is a model Pr of D^p w.r.t. $\mathcal{IC}$ such that $p(t_1 \wedge \cdots \wedge t_n)$ w.r.t. Pr is equal to $\sum_{\substack{\alpha \in B^{n+m} \mid \\ \forall i \in [1..n] \alpha[i] = true}} x_\alpha$, and vice versa.

Given a solution σ of $S(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$, for each $\alpha \in B^{n+m}$ we denote with σ_α the value assumed by the variable x_α in σ; moreover, for each $(t_i, t_j) \in EP$ we denote with σ_{t_i,t_j} the value assumed by the variable x_{t_i,t_j} in σ.

For each $(t_i, t_j) \in EP$, we denote with $D^p_{t_i,t_j}$ the maximal subset of D^p which contains only t_i, t_j, and the tuples along the path connecting t_i and t_j.

From Proposition 2, the fact that, for each $(t_i, t_j) \in EP$, the value σ_{t_i,t_j} is such that $p^{min}(t_i \wedge t_j) \leq \sigma_{t_i,t_j} \leq p^{min}(t_i \wedge t_j)$, implies that there is at least a model Pr_{t_i,t_j} of $D^p_{t_i,t_j}$ w.r.t. $\mathcal{IC}$ such that $p(t_i \wedge t_j)$ w.r.t. Pr_{t_i,t_j} is equal to σ_{t_i,t_j}. For each $(t_i, t_j) \in EP$, we consider a model Pr_{t_i,t_j} of $D^p_{t_i,t_j}$ w.r.t. $\mathcal{IC}$ such that $p(t_i \wedge t_j)$ w.r.t. Pr_{t_i,t_j} is equal to σ_{t_i,t_j}. Moreover, for each possible world $w \in pwd(D^p_{t_i,t_j})$, we define the relative weight of w (and denote it by $wr(w)$) as:

$$wr(w) = \begin{cases} \dfrac{Pr_{t_i,t_j}(w)}{\sum_{w' \in D^p_{t_i,t_j} \wedge t_i \in w' \wedge t_j \in w'} Pr_{t_i,t_j}(w')} & \text{if } t_i \in w \wedge t_j \in w \\ \dfrac{Pr_{t_i,t_j}(w)}{\sum_{w' \in D^p_{t_i,t_j} \wedge, t_i \in w' \wedge t_j \notin w'} Pr_{t_i,t_j}(w')} & \text{if } t_i \in w \wedge t_j \notin w \\ \dfrac{Pr_{t_i,t_j}(w)}{\sum_{w' \in D^p_{t_i,t_j} \wedge t_i \notin w' \wedge t_j \in w'} Pr_{t_i,t_j}(w')} & \text{if } t_i \notin w \wedge t_j \in w \\ \dfrac{Pr_{t_i,t_j}(w)}{\sum_{w' \in D^p_{t_i,t_j} \wedge t_i \notin w' \wedge t_j \notin w'} Pr_{t_i,t_j}(w')} & \text{if } t_i \notin w \wedge t_j \notin w \end{cases}$$

It is easy to see that, for each possible world $w \in pwd(D^p)$, there is for each pair $(t_i, t_j) \in EP$ a possible world $w_{t_i,t_j} \in pwd(D^p_{t_i,t_j})$ such that $w = \bigcup_{(t_i,t_j)\in EP} w_{t_i,t_j}$, and vice versa.

We consider the interpretation Pr of D^p defined as follows. For each possible world $w \in pwd(D^p)$, we consider the possible worlds w_{t_i,t_j} such that $w = \bigcup_{(t_i,t_j)\in EP} w_{t_i,t_j}$ and define the interpretation Pr of D^p as:

$$Pr(w) = \sigma_\alpha \prod_{(t_i,t_j)\in EP} wr(w_{t_i,t_j}),$$

where α is the tuple in B^{n+m} which agrees with w on the presence/absence of $t_1, \cdots, t_{n+m}$ (i.e., $\forall i \in [1..n+m]\ \alpha[i] = \text{true}$ (resp. *false*) iff $t_i \in w$ (resp. $t_i \notin w$)). It is easy to see that Pr is a model for D^p w.r.t. $\mathcal{IC}$. Specifically, the following conditions hold:

- Pr assigns probability 0 to every possible world w not satisfying $\mathcal{IC}$. This can be proved reasoning by contradiction. Assume that $Pr(w) > 0$ and w does not satisfy $\mathcal{IC}$. Consider the possible worlds w_{t_i,t_j} such that

 $$w = \bigcup_{(t_i,t_j)\in EP} w_{t_i,t_j}.$$

 Since $Pr(w) > 0$, for each $(t_i, t_j) \in EP$ it holds that

 $$Pr_{t_i,t_j}(w_{t_i,t_j}) > 0.$$

 Hence, since Pr_{t_i,t_j} is a model of $D^p_{t_i,t_j}$, then w_{t_i,t_j} contains no pair of tuples t', t'' connected by an edge in $HG(D^p, \mathcal{IC})$. Therefore, w contains no pair of tuples t', t'' connected by an edge in $HG(D^p, \mathcal{IC})$, thus contradicting that w does not satisfy $\mathcal{IC}$.

- For each tuple $t \in D^p$, $p(t) = \sum_{w\in pwd(D^p)\wedge t\in w} Pr(w)$. This follows from the fact that, given a tuple $t \in D^p$, and such that t belongs to a chain whose ends are the tuples t_i, t_j, the probability of a tuple t is given by

 $$\sum_{w_{t_i,t_j}\in pwd(D^p_{t_i,t_j})\ s.t.\ t\in w_{t_i,t_j}} Pr_{t_i,t_j}(w_{t_i,t_j}).$$

 The latter is equal to $\sum_{w\in pwd(D^p) s.t. t\in w} Pr(w)$, since for each $w_{t_i,t_j} \in pwd(D^p_{t_i,t_j})$ it holds that

 $$\sum_{w\in pwd(D^p)\ s.t.\ w_{t_i,t_j}\subseteq w} Pr(w) = Pr_{t_i,t_j}(w_{t_i,t_j}).$$

Therefore, the interpretation Pr is a model for D^p w.r.t. $\mathcal{IC}$, and the probability assigned to $t_1 \wedge \cdots, t_n$ by Pr is equal to $\sum_{\substack{\alpha\in B^{n+m} \\ \wedge\forall i\in[1..n]\,\alpha[i]=\text{true}}} \sigma_\alpha$. Hence, it is easy to see that $p^{min}(t_1 \wedge \cdots \wedge t_n)$ (resp. $p^{max}(t_1 \wedge \cdots \wedge t_n)$) is the optimal solution of $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ and can be computed in polynomial time w.r.t. the size of D^p, which completes the proof. □

Theorem 13. For projection-free queries, QA is in $PTIME$ if $HG(D^p, \mathcal{IC})$ is a simple graph.

Proof. Let $\vec{t}$ be an answer of the projection-free query Q posed on the deterministic version of D^p. The minimum and maximum probabilities p^{min} and p^{max} of $\vec{t}$ as answer of Q over D^p can be determined as follows. Let $T = \{t_1, \ldots, t_n\}$ be the set of tuples in D^p such that $Q(\vec{t}) = t_1 \wedge \cdots \wedge t_n$. T can be partitioned into the sets $T_1, \ldots, T_k$, such that:
1) k is the number of distinct (maximal) connected components of $HG(D^p, \mathcal{IC})$, each of which contains at least one tuple in T;
2) for each $i \in [1..k]$, T_i contains the tuples of T belonging to the i-th maximal connected component of $HG(D^p, \mathcal{IC})$ among those mentioned in 1).
Let $\vec{t_i}$ be the conjunction of the tuples belonging to the partition T_i of T. Since every maximal connected component of $HG(D^p, \mathcal{IC})$ is either a clique or a tree, lemmas 5 and 6 ensure that $p^{min}(\vec{t_i})$ and $p^{max}(\vec{t_i})$ can be computed in polynomial time w.r.t. the size of D^p. As distinct tuples $\vec{t_i}$ and $\vec{t_j}$, with $i, j \in [1..k]$, belong to distinct maximal connected components of $HG(D^p, \mathcal{IC})$, they can be viewed as events among which no correlation is known. Hence, $p^{min}(\vec{t})$ (resp., $p^{max}(\vec{t})$) can be determined by applying Fact 2 to the events $\vec{t_1}, \ldots \vec{t_k}$, with the probability of $\vec{t_i}$ equal to $p(\vec{t_i}) = p^{min}(\vec{t_i})$ (resp., $p(\vec{t_i}) = p^{max}(\vec{t_i})$), for each $i \in [1..k]$. □

Appendix A.7. Extending tractable cases of query evaluation

As discussed in the core of the paper (Section 6), our tractability result on query evaluation can be extended to the cases that: *i*) tuples are associated with ranges of probabilities, instead of exact probability values; *ii*) denial constraints are probabilistic. We here give a hint on how the proof of Lemma 6 can be extended to these cases (Lemma 6 states that projection-free queries can be evaluated in *PTIME* if the conflict hypergraph is a tree, and is the core of the proof of Theorem 13).

As regards extension *ii*), it is easy to see that, as shown for CC, any instance I of the query evaluation problem in the presence of probabilistic constraints is equivalent to an instance I' of QA, where the conflict hypergraph H' of I' is obtained by augmenting each hyperedge of the conflict hypergraph H of I with an ear. The point is that, even if H is a tree, this reduction makes H' contain hyperedges with more than two nodes, thus H' is no more a tree. However, H' is a hypertree of a particular form: for any pairs of intersecting edges, their intersection consists of a unique node, which is a node inherited from H (the new nodes of H' are all ears). This implies that the minimum and maximum probabilities p^{min} and p^{max} of an answer can be still computed as solutions of the two variants of the optimization problem $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ introduced in the proof of Lemma 6. The fact that $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ can be still written and solved in polynomial time derives from the fact that the values $p^{min}(t_i \wedge t_j)$ and $p^{max}(t_i \wedge t_j)$ occurring in the inequalities (B) can be still evaluated in polynomial time, by observing that both $p^{min}(t_i \wedge t_j)$ and $p^{max}(t_i \wedge t_j)$ can be obtained by exploiting Lemma 3 for the minimum probability value, and an analogous result for the maximum probability value. Observe that this reasoning does not work (as is) for general hypertrees, as in this case we are not assured that the tuples composing the answer are in intersections between distinct pairs of hyperedges.

As regards extension *i*), the minimum and maximum probabilities p^{min} and p^{max} of an answer can be computed as solutions of the two variants of the optimization problem $LP(t_1 \wedge \cdots \wedge t_n, \mathcal{D}^p, \mathcal{IC}, D^p)$ with the following changes:
1) equalities (C) are replaced with pairs of inequalities imposing that, for each t_i, its probability ranges between the minimum and maximum marginal probabilities of the range associated with t_i in the PDB;
2) the values $p^{min}(t_i \wedge t_j)$ and $p^{max}(t_i \wedge t_j)$ occurring in the inequalities (B) are evaluated by considering the minimum probabilities for the tuples along the path connecting t_i and t_j in the conflict tree. Moreover, when evaluating $p^{min}(t_i \wedge t_j)$, the minimum marginal probabilities for t_i and t_j are taken into account, while, for $p^{max}(t_i \wedge t_j)$, we have to consider their maximum probabilities. Therein, the maximum probability of a tuple t is the minimum between the upper bound of the probability range of t, and the maximum probability value that t can have according to the conflict tree (this value is entailed by the tuples connected to t by direct edges: as implied by Theorem 2, the sum of the probabilities of two tuples connected through an edge must be less than or equal to 1). Intuitively enough, we consider the minimum probabilities for the intermediate tuples between t_i and t_j as this allows the greatest degree of freedom in distributing t_i and t_j in the probability space.